

Over-Determination and Act-Consequentialism

Magnus Jedenheim Edling



UPPSALA
UNIVERSITET

Dissertation presented at Uppsala University to be publicly examined in Geijersalen, Thunbergsvägen 3P, Uppsala, Thursday, 4 May 2017 at 10:00 for the degree of Doctor of Philosophy. The examination will be conducted in English. Faculty examiner: Professor David Boonin (University of Colorado Boulder).

Abstract

Jedenheim Edling, M. 2017. *Over-Determination and Act-Consequentialism*. 112 pp. Uppsala: Department of Philosophy, Uppsala University. ISBN 978-91-506-2630-8.

This dissertation is a discussion of the challenge that cases of over-determination pose to Act-Consequentialism. Although there are many realistic examples of such cases – for example, pollution, overfishing, or the election of an inappropriate politician – I consider structurally purer examples, one of which I call “Case One.” Suppose that you and I independently shoot and kill a third person called “Victim.” Our bullets arrive at the same time and each shot would have killed Victim by itself. Finally, Victim would not have been killed, if neither of us had pulled the trigger. According to the Standard Version of Act-Consequentialism, an action is wrong if and only if it has an alternative whose consequences would be intrinsically better. Case One challenges the Standard Version because there does not seem to be such an alternative to my action: Victim would have died by your shot if I had not shot him, and similar remarks apply to your action.

The dissertation is structured as follows. After Chapter One, which briefly introduces the main issues of the dissertation, I turn to Chapter Two – “Preliminaries” – where I outline the Standard Version and highlight the main characteristics of over-determination cases. These cases are divided into cases of redundant difference making and cases of redundant causation. Cases of redundant causation are subdivided further into cases of causal over-determination and pre-emption. I make an important stipulation in this chapter. I say that our actions in Case One and similar cases are “redundant negative difference makers.”

In Chapter Three – “Replies” – I consider whether the proponent of Act-Consequentialism might question the intuition that you and I, respectively, act wrongly in Case One. The proponent might accept that we have this intuition but explain it away, or she might deny that we have the intuition and instead point to something else that is wrong in this kind of case. For example, she could suggest that although neither you nor I act wrongly individually, we act wrongly together. I argue that these replies are problematic. For instance, explaining away this intuition might also force us to explain away intuitions that support the Standard Version.

In Chapter Four – “Causal Consequences” – I discuss an alternative version of Act-Consequentialism that might seem to fare better. The Standard Version interprets the term “outcome of an action” as referring to the entire possible world that would obtain, if the action were performed. The version I have in mind, the “Causal Consequences Version of Act-Consequentialism,” understands “outcome of an action” as instead referring to the causal consequences of the action. It seems clear that you and I, respectively, cause the state of affairs that Victim dies in Case One. However, I show that the Causal Consequences Version has a number of unattractive implications.

In Chapter Five – “the Non-Standard Version” – I suggest another alternative version of Act-Consequentialism. This version – the Non-Standard Version – implies that you and I act wrongly in Case One. Roughly, the Non-Standard Version says that an action is wrong if and only if it has an alternative whose consequences would be intrinsically better, or is a redundant negative difference maker. The Non-Standard Version is similar to a principle suggested by Derek Parfit. However, I shall argue that the Non-Standard Version is preferable to Parfit’s principle.

In Chapter Six – “Further Cases” – I discuss a number of cases that challenge the Non-Standard Version. For example, what would the Non-Standard Version imply in a case very similar to Case One but where I would have killed another person, if I had not shot Victim? I argue that the Non-Standard Version handles this and other problematic cases, and that it is therefore a plausible alternative to the Standard Version.

Keywords: Over-Determination, Pre-Emption, Act-Consequentialism

Magnus Jedenheim Edling, Department of Philosophy, Ethics and Social Philosophy, Box 627, Uppsala University, SE-75126 Uppsala, Sweden.

© Magnus Jedenheim Edling 2017

ISBN 978-91-506-2630-8

urn:nbn:se:uu:diva-318643 (<http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-318643>)

Contents

Chapter One, Introduction	9
1. Introduction	9
2. A Brief Description of the Challenge.....	10
3. A Brief Description of Possible Replies	12
4. Upshot	13
3. Plan of the Thesis	14
Chapter Two, Preliminaries	16
1. Introduction	16
2. The Standard Version of Act-Consequentialism.....	17
2.1 Intrinsic Value	18
2.2 Outcomes	18
2.3 Actions.....	19
2.4 Alternatives.....	21
2.5 Negative Difference Makers.....	22
3. Redundant Negative Difference Making.....	22
3.1 Redundant Negative Difference Makers	23
3.2 Cases Involving More Than Two Actions.....	26
3.3 A Case Involving Dependence	28
4. Redundant Causation.....	29
4.1 Redundant Causes.....	30
4.2 Case Three and Case Four	32
4.3 Case Five	33
4.4 Over-Determination and Pre-Emption.....	34
Chapter Three, Replies.....	36
1. Introduction	36
2. The Challenge	37
3. Conceding the Intuition.....	38
3.1 Denying the Evidential Value of the Intuition.....	39
3.2. Partners in Crime	40
4. Denying the Intuition	41
5. Acting Wrongly Together	42
5.1 The Group Action Is Wrong.....	42
5.2 Tännsjö's Reply	43
5.3 Jackson's Reply	45
5.4 An Appeal to Pluralities of Actions.....	46

6. An Appeal to Subjective Moral Principles.....	47
6.1 Subjective Moral Principles.....	48
6.2 Subjective Moral Principles and Case One and Case Two	50
7. An Appeal to Uncooperativeness.....	51
Chapter Four, Causal Consequences.....	52
1. Introduction.....	52
2. The Causal Consequence Version.....	53
2.1 Case One and Case Two.....	54
3. A Difficulty for The Causal Consequence Version	55
3.1 The Simple Counterfactual Account of Causation	56
3.2 Case One.....	57
4. A Helping Hand	59
4.1 An Objection to the Simple Account.....	59
4.2 Causation	60
4.3 Causation and Case Two	61
4.4 Causation and the “Ordinary” Case.....	62
4.5 Causation and Redundant Causes.....	63
5. Rejecting the Causal Consequence Version.....	64
5.1 The Causal Consequence Version and Ordinary Cases.....	64
5.2 The Intrinsic Value of Death	67
Chapter Five, The Non-Standard Version	69
1. Introduction	69
2. The Non-Standard Version.....	70
2.1 The Non-Standard Version and Case One.....	70
2.2 Case Three and Case Four	73
2.3 Moral Dilemmas	74
2.4 Innocent Agents.....	77
2.5 Two Worries	78
3. Other Deontic Categories.....	79
3.1 Rightness and Obligatoriness	79
3.2 Redundant Positive Difference Making	80
3.3 The Symmetry Objection.....	81
3.4 A Mixed Case	82
3.5 A Further Case.....	84
4. The Parfittian Principle	85
4.1 Introducing the Parfittian Principle, and Case One	86
4.2 Case Three and Case Four	87
4.3 Case Seven and Case Eight	89
4.4 The Non-Standard Version vs. the Parfittian Principle	90

Chapter Six, Further Cases	92
1. Introduction	92
2. Pre-Emption	92
2.1 Case Two	93
2.2 A Worry	94
2.3 When Your Action Occurs Before Mine	95
2.4 When Your Action Occurs After Mine	97
2.5 When Your Action Does Not Occur	98
3. Cases with Better or Worse Alternatives	100
3.1 Cases Where All Agents Have Better or Worse Alternatives	101
3.2 Cases Where Some Agents Have Better or Worse Alternatives	102
3.3 Modifying the Non-Standard Version	104
3.4 The Non-Standard Version is “Exclusively Act-Oriented”	106
3.5 A Further Objection	108
4. Cases with Natural Causes	109
Bibliography	111

Acknowledgements

I want to thank everyone who has helped me in the course of writing this thesis. I am extremely grateful to my supervisors, Folke Tersman and Jens Johansson, for their encouragement and support. I have benefitted greatly from their generous comments – both general and detailed – on my work.

I also want to thank Krister Bykvist and Erik Carlson. Krister Bykvist provided very helpful comments on the whole draft of my thesis in a seminar at Uppsala University, April 2016. Chapter Four of my thesis originated from his comments. Erik Carlson has given me invaluable comments on my thesis over the years. Some of these comments were counterexamples to the principles I tried to construe, which was frustrating at the time, but has benefitted my thesis enormously. Tobias Wilsch generously provided helpful comments on a previous draft of Chapter Two.

I have present all chapters of my thesis at the higher seminar in Practical Philosophy, Uppsala University. I have presented part of my thesis at Filosofidagarna in Gothenburg, June 2011, and at Filosofidagarna in Linköping, June 2015. I am grateful to all participants at these seminars, especially, Per Algander, Emil Andersson, Katharina Berndt-Rasmussen, Christian List, Karin Enflo, Karl Ekendahl, Per Ericson, Björn Eriksson, Patricia Mindus, Victor Moberger, Jessica Pepp, Karl Pettersson, Olle Risberg, Andrew Reisner, and Simon Rosenqvist.

I also want to than Gästrike-Hälsinge Nation and Helge Ax:son Johnsons stiftelse for their generous financial support. I am particularly grateful to Gästrike-Hälsinge Nation for awarding me the Göransson-Sandviken Research Scholarship.

Finally, I am also very grateful to the Department of Philosophy, Uppsala University, and all the people working at the department, researchers as well as administrative personnel. It is a great place to work in.

Chapter One, Introduction

1. Introduction

There are numerous cases where the effect of what a number of persons do is negative, but where the effect of what each person does is not. For example, think of cases of overfishing where the catch each fisherman takes has little or no effect on the stock of fish but where what they together take out seriously depletes the stock or even leads to a complete collapse. Or think of cases of pollution where the emissions of each workplace or each household have little or no effect but where their combined emissions have serious environmental effects such as, for example, the depletion of the ozone layer, algal bloom, acidification of lakes, etc. Or think, finally, of elections where each abstention does not make any difference but where all abstentions together allow for the election of an inappropriate political party or an inappropriate politician, or worse. There is no question that these cases, and cases like these, constitute a real danger. Indeed, some people think that they may even threaten the survival of mankind.

There are several questions we might ask as regards these cases. For example, consider the fishermen in an overfishing case such as the ones described above. Suppose their fishing led to a complete collapse of the stock and that they could not influence the behaviour of one another. We might then ask whether they acted *rationally*. It might seem odd to say that they did. After all, they undermined their livelihood through their own behaviour. On the other hand, there are reasons in favour of the view that each fisherman did act rationally: The others were going to fish anyway; so each fisherman might just as well try to catch as much fish as possible before the stock collapses.

We might also ask whether they *caused* the collapse. On the one hand, it would be odd to claim that the fishermen did not do so. The stock did collapse; and if the fishermen did not do it, then who did? However, there are also reasons in favour of saying that the actions of the fishermen did not cause the collapse. Causes are plausibly thought of as being necessary for their effects, and no individual action was necessary in this case. Finally, we might ask whether the fishermen acted *wrongly*. It seems weird to say that they did not. Suppose that all they had to eat was fish. In that case, they would all die, and their kids would too. But there are reasons in favour of claiming that no individual fisherman acted wrongly. The moral status of an action – for example, whether an action is wrong – is plausibly determined by the well-being effects

of the action – at least consequentialists think so – and no individual action did have any well-being effects.

In this thesis I am concerned with the third question. More specifically, I am concerned with the challenge these cases pose to what I call the “Standard Version of Act-Consequentialism”. To explain this challenge better, I need to say a bit more about the Standard Version of Consequentialism and these cases.

2. A Brief Description of the Challenge

According to the Standard Version of Act-Consequentialism, an action *a* is wrong, if and only if, there is an alternative to *a* whose outcome would be intrinsically better than that of *a*.

I am going to discuss the Standard Version at length in this thesis. At this stage, it is sufficient to note that it has the correct implication in a typical case of wrongdoing. Suppose that Gertrude kills Victim who has a life well worth living in front of him; and that Gertrude could instead have done many other things, as for example, visit her sickly aunt Dottie. In that case, the Standard Version implies that Gertrude acts wrongly. She has an alternative whose outcome would be intrinsically better: It would have been better if she had visited her aunt Dottie, for example.

Although there are many realistic cases where the effect of what a number of persons do is negative but where the effect of what each person does is not – as the case of overfishing mentioned above, for example – I am going to discuss simpler but structurally analogous cases in this thesis. There are particularly two cases that are important, which I call “Case One” and “Case Two.” These two cases are – among other things – important because the other cases I shall discuss in this thesis are based on them.

Consider first Case One:

Case One

You and I shoot and kill Victim. Our bullets arrive at the same time and each shot would have killed Victim by itself. Victim would not have been killed, if neither of us had pulled the trigger. Apart from shooting Victim, you and I have only one alternative each: each could also have abstained from shooting Victim. Furthermore, each of us would have performed his or her alternative – i.e., each of us would have abstained from shooting Victim – if he or she had not shot him. We act independently of one another, Victim’s level of well-being is not affected by either of the shots taken by itself (e.g., Victim does not experience any pain), and no one else is affected. Finally, Victim would have gone on to live a life well worth living, if neither you nor I had shot him.

I am not only going to discuss the Standard Version at length in this thesis; I shall also discuss Case One and cases like Case One at length. However, now it is enough to note that neither you nor I act wrongly in Case One, according to the Standard Version. There is no alternative to what I do whose outcome would be intrinsically better: If I had not shot Victim you would still have done so. Parallel things are true of what you do as well. Consequently, neither of acts wrongly, according to the Standard Version. However, this is counter-intuitive: Victim dies as a consequence of what you and I do. Yet, neither you nor I act wrongly.

Before I turn to Case Two there are a few things I would like to note as regards my brief discussion of Case One. First, I concluded that I do not have an alternative whose outcome would be intrinsically better *because* you would have killed Victim in any event, if I had not shot him. This motivation is not – strictly speaking – sufficient. Other things must be true as well. For example, it must be true that my shot does not cause Victim’s level of well-being to fall, say, by causing him severe pain (which it does not, according to the description of Case One). Otherwise, I would have an alternative whose outcome is intrinsically better (although you would in any event have killed Victim). Victim would under those circumstances at least have been spared the pain, if I had refrained from shooting him. I am nevertheless not always going to present all premises I rely on when concluding that an action is wrong, for example. It is evident that I rely on them from the description of the cases.

Second, I also concluded that it is counterintuitive that neither you nor I act wrongly in Case One, according to the Standard Version *because* Victim dies as a consequence of what you and I do. (I shall offer further reasons why this might seem counterintuitive in later chapters of this thesis.) However, this motivation is also not – strictly speaking – sufficient. If Victim would have experienced an unimaginable pain for the rest of a his (long) life, if you and I had not killed him (which he would not, according to the description of Case One), it would probably not be counterintuitive that neither you nor I act wrongly by killing him. Again, I am not always going to present all premises I rely on since it is evident that I rely on them.

Consider now this case:

Case Two

You and I shoot one shot each at Victim. My bullet arrives first and kills Victim. Your bullet arrives very shortly after mine but does not kill Victim since my bullet has already done that. Your bullet would however have killed Victim, if my bullet had not done so first. Victim would not have been killed, if neither of us had pulled the trigger. Apart from shooting Victim, you and I have only one alternative each: each could also have abstained from shooting Victim. Furthermore, each of us would have performed his or her alternative – i.e., each of us would have abstained from

shooting Victim – if he or she had not shot him. We act independently of one another, Victim’s level of well-being is not affected by my shot (e.g., Victim does not experience any pain), and no one else is affected. Finally, Victim would have gone on to live a life well worth living, if neither you nor I had shot him, and the extra time Victim would have lived, if I had not shot him would neither have been intrinsically good nor intrinsically bad for him.

This case is a bit different from Case One: I but not you cause Victim’s death here, for example. However, there is no need to dwell on the differences between Case One and Case Two for the time being. (As I said, I am going to discuss these cases and cases like them thoroughly in this thesis.) At this point, it is enough that we note that neither you nor I act wrongly in this case either, according to the Standard Version. Again, there is no alternative to what I do whose consequences would be better: If I had not shot Victim you would still have done so, and similar remarks are true of you too. However, this is counterintuitive in this case as well. (It might be thought that it is not counterintuitive that *you* do not act wrongly in Case Two since you do not actually kill Victim in this case. As we shall see, I shall assume that a moral principle *may* imply that you act wrongly, however.)

Note that I concluded in this case too that you do not have an alternative whose outcome would be intrinsically better without explicitly mentioning all the premises I rely on. For example, if the extra time you would have lived if I had not shot you would have been intrinsically bad for you because say, you would have experienced a lot of pain (which he would not, according to the description of Case One), I would have an alternative whose outcome would be intrinsically better. However, as I said above, I am not always going to make all premises I rely on explicit.

3. A Brief Description of Possible Replies

The act-consequentialist has two options here. On the one hand, she could hold on to the Standard Version. In that case, she would have to insist that the Standard Version has the correct implications in Case One and in Case Two. That is to say, she would have to insist that neither you nor I do in fact act wrongly in these cases. On the other hand, she could reject the Standard Version and adopt some other version of Act-Consequentialism that implies that you and I do act wrongly in these cases (or that at least implies that both do so in Case One, and that I do so in Case Two).

If she chooses the first option there seem to be two available strategies: First, she may concede that the implications of the Standard Version in Case One and in Case Two are counterintuitive, and try to find a way to deal with that. For example, she might try to explain away this intuition. Second, she

may deny that the implications of the Standard Version in Case One and in Case Two are counterintuitive. She could claim that we have the less specific intuition that something is not as it should be in these cases rather than the intuition that you and I, respectively, act wrongly. She could then come up with a suggestion of what it is that is not as it should be in these cases. As we shall see, there are primarily three suggestions of what that is in the literature: that you and I together act wrongly, that you and I respectively act subjectively wrongly, and that you and I are uncooperative.

If she chooses the second option, there are bound to be many alternatives. However, I think that there are two that stand out: a view that I call the “Causal Consequences Version of Act-Consequentialism” and a view that I call the “Non-Standard Version of Act-Consequentialism.”

As we shall see, the Standard Version interprets the term “outcome of an action” as referring to the possible world that would be the case, if the action were performed. The Causal Consequences Version instead interprets the same expression as referring to the complete causal consequences of the action that would be the case, if the action were performed. This might initially seem to be a sensible move. Is it not clear that you and I, respectively, cause Victim’s death in Case One and that I cause Victim’s death in Case Two?

The Non-Standard Version does not claim, as the Standard Version does, that an action is wrong only in virtue of having an alternative whose outcome is intrinsically better. Rather, it claims that an action is wrong in virtue of having an alternative whose outcome is intrinsically better or being what I call a “redundant negative difference maker”. I am going to have quite a lot to say about redundant negative difference makers in this thesis. For now, it is sufficient – in order to get some understanding of what sort of things redundant negative difference makers are – to note two things. First, it is fitting to call an action that is wrong, according to the Standard Version a “negative difference maker” since such an action makes a difference for the worse. Gertrude (above) makes a difference for the worse by killing Victim, for example. Second, our actions in Case One and Case Two, respectively, are such that each would have been a negative difference maker if it were not for the other. It thus seems fitting to call them “redundant negative difference makers.”

I am going to consider both these options in this thesis.

4. Upshot

I am inclined to favour the Non-Standard Version. It has (I believe) the intuitively correct implications in Case One and Case Two, and also in many other similar cases. I do not think that the Causal Consequence Version is particularly promising. As we shall see, it has some pretty implausible implications. For example, it has the intuitively correct implications in Case One and Case

Two, only if we assume that Victim's death is intrinsically bad for him. However, I grant that both strategies (particularly the second one) of the first option also have something going for them. The reason I favour the Non-Standard Version is primarily that I cannot rid myself of the impression that you and I, respectively, act wrongly in Case One and Case Two (and in many other comparable cases).

I recognise that most people I have talked to that are also familiar with the ins and outs of this discussion favour the first option. Among them, there are those that agree that the Non-Standard Version has the intuitively correct implications in Case One and Case Two, and in other similar cases – and consequently that the Standard Version does not have the intuitively correct implications in these cases – but think that it is more important that a moral theory satisfies some other theoretical standard. For example, some think that I undervalue theoretical simplicity. (As we shall see, the Non-Standard Version is quite complex.) There are also those who do not think that the Non-Standard Version has the intuitively correct implications in Case One and Case Two, or in other comparable cases (but that the Standard Version does have the intuitively correct implications in these cases). However, even if they are right about this, I think that both strategies of the first option are sufficiently problematic to motivate an inquiry into whether the second option mentioned above (i.e., to reject the Standard Version and adopt some other version of Act-Consequentialism) is tenable, or so I shall argue.

To be clear, I do not take myself to show in this thesis that the Non-Standard Version is preferable to the Standard Version. Rather, what I want to suggest is that the Non-Standard Version is a moral theory worthy of consideration.

3. Plan of the Thesis

In Chapter Two, I shall discuss both Case One and Case Two, and the Standard Version more carefully in order to be able to better address the challenge posed by these cases. I shall stipulate that our actions in Case One and Case Two are *redundant negative difference makers*. As I said in section 3, a proponent of Act-Consequentialism might give up on the Standard Version and instead embrace the Non-Standard Version, which says (among other things) that an action is wrong in virtue of having an alternative whose outcome is intrinsically better *or* being a redundant negative difference maker. (This is also what I do in Chapter Five and Chapter Six.)

In Chapter Two, I shall moreover say that our actions in Case One and Case Two also are *redundant causes* of Victim's death. As I also said in section 3, a proponent of Act-Consequentialism might abandon the Standard Version and instead adopt the Causal Consequences Version. Intuitively, Victim's death – in Case One, for example – is a causal consequence of your action and

my action, respectively. A proponent of the Causal Consequence Version might then argue that the fact that our actions are redundant causes of Victim's death supports this intuition. For example, she might suggest that an action is a cause, if it among other things is a redundant cause. (I shall discuss this view and the Causal Consequences Version (which I reject) in Chapter Four.)

In Chapter Three, I consider some replies that the proponent of the Standard Version might make, namely the ones I described in section 2 above.¹

In Chapter Four, I shall discuss the Causal Consequence Version. I argue that we should reject it because it has some unattractive implications.

In Chapter Five, I am going to defend the Non-Standard Version. I show that it has the intuitively correct implications in cases like Case One and other similar cases. I note that there are cases where it is not evident that the Non-Standard Version has the intuitively correct implications. However, I also suggest that it is not evident that it does not have the intuitively correct implications in these cases.

Finally, in Chapter Six, I argue that the Non-Standard Version has the correct implication in Case Two, and other similar cases. I note here too that there are cases where it is not clear that the Non-Standard Version has the intuitively correct implications and also that it is not clear that it does not have the intuitively correct implications in these cases. I also argue that the Non-Standard Version has the correct implication in a number of other cases relevant in assessing the Non-Standard Version. For example, there are cases like Case One but where the agents have alternatives whose outcome is either intrinsically better or intrinsically worse.

¹ Of course, a consequentialist might abandon the Standard Version and instead adopt Rule-Consequentialism or Consequentialist Generalisation to deal with Case One and Case Two. However, in this thesis I am concerned with what a proponent of the Standard Version might say in the face of these cases. In Chapter Six (section 3.4), I shall also point out that there are cases very similar to Case One and Case Two where Rule-Consequentialism and Consequentialist Generalisation seem to have counterintuitive implications.

Chapter Two, Preliminaries

1. Introduction

In the previous chapter I noted that Case One and Case Two (i.e., the cases where you and I shoot and kill a third person called “Victim”) challenge the Standard Version of Act-Consequentialism. In this chapter I am going to describe both Case One and Case Two, and the Standard Version a bit more carefully. By doing so, we will be better equipped to address the challenge posed by Case One and Case Two. As we shall see throughout this thesis, there are many reasons why it is a good idea to study Case One and Case Two, and the Standard Version closer. I mention two, for the time being.

First, in this chapter I shall stipulate that our actions in Case One and Case Two are *redundant negative difference makers*. A proponent of Act-Consequentialism might give up on the Standard Version and instead embrace the what I call the “Non-Standard Version,” which roughly suggests that an action is wrong in virtue of having an alternative whose outcome is intrinsically better *or* being a redundant negative difference maker. In fact, I shall defend this view in Chapter Five and Chapter Six.

Second, in this chapter I shall say that our actions in Case One and Case Two are also *redundant causes* of Victim’s death. A proponent of Act-Consequentialism might abandon the Standard Version and instead adopt what I call the “Causal Consequences Version of Act-Consequentialism.” Intuitively, Victim’s death – in Case One, for example – is, I believe, a causal consequence of your action and my action, respectively. A proponent of the Causal Consequence Version might argue that the fact that our actions are redundant causes of Victim’s death supports this intuition. For example, she might want to suggest that an action is a cause, if among other things it is a redundant cause. I shall discuss this view and the Causal Consequences Version (which I reject) in Chapter Four.

This section is organised as follows. In section 2, I discuss and identify the main features of the Standard Version. In section 3, I say under what circumstances actions are redundant negative difference makers. I go on to show that our actions in Case One and Case Two are redundant negative difference makers. I also discuss a number of other cases here. For example, I discuss cases that differ from Case One and Case Two just in that they involve more people. I show that the actions of those involved in these cases are redundant

negative difference makers too. In section 4, I shall say under what circumstances an action is a redundant cause of another event. I then show that the actions in Case One and Case Two – and also in the cases that involve more people – are redundant causes of Victim’s death. I also distinguish various cases of redundant causation here. They fall into two large groups: cases of over-determination (where Case One belongs) and cases of pre-emption (where Case Two belongs).

2. The Standard Version of Act-Consequentialism

The Standard Version of Act-Consequentialism is of course a version of Act-Consequentialism. I shall therefore begin by characterising the latter view. Act-Consequentialism is broadly speaking the view that each person – given the circumstances she finds herself in – ought to make consequences as good as possible. Act-Consequentialism may more specifically be formulated as follows:

Act-Consequentialism

An action *a* is *wrong*, if and only if, there is an alternative to *a* whose outcome would be intrinsically better than that of *a*.

An action *a* is *obligatory*, if and only if, the outcome of *a* is intrinsically better than the outcome of any of its alternatives would be.

An action *a* is *right*, if and only if, there is no alternative to *a* whose outcome would be intrinsically better than that of *a*.

Act-Consequentialism may be spelled out in several different ways. For example, there are several views on what an action is. There are also several views on what the outcome of an action is, and so on. The Standard Version makes two claims that set it apart from other versions of Act-Consequentialism. First, it says that the intrinsic value of an outcome is determined by *welfare*. Second, it says that the outcome of an action is the *total* way things would have been if the action were performed. I shall discuss the first of these claims in section 2.1 and the second in section 2.2. Furthermore, I shall discuss what actions are in section 2.3 and what the alternatives to an action are in section 2.4. The Standard Version does not take a particular position on either of these two issues. Nevertheless, we need to discuss them too. Finally, in section 2.5, I shall highlight an important feature of the Standard Version.

Before I turn to these matters, I want to point out that Act-Consequentialism is a maximising view. It says that the agent should perform the action available to her in the circumstances whose outcome has the highest intrinsic

value. Many I suspect think that it would be implausible to hold the view that the moral status of an action is exclusively determined by the relative intrinsic value of its outcome without also holding the view that intrinsic value should be maximised. There are consequentialists that hold the former view without holding the latter, however. They do so to respond, for example, to the objection that consequentialism is too demanding. For instance, Act-Consequentialism implies that you should sacrifice your life in order to save the lives of two strangers, if this maximises intrinsic value.

2.1 Intrinsic Value

The Standard Version claims that it is *welfare* or *well-being* (and only welfare or well-being) that determines the intrinsic value of an outcome. Welfare is a prudential value: it is, that is to say, concerned with how life goes for the person whose life it is. There are of course various theories of what the determinants of welfare are – for example, Hedonism and Preferentialism – but we do not need to take a stand here.² Our discussion does not depend on this issue.

2.2 Outcomes

According to the Standard Version, the outcome of an action is the *total* way things would be if the action were performed. More precisely, according to the Standard Version, the outcome of an action is the *entire* possible world that would be the case, if the action were performed. Some think that we should instead only count the future of the possible world that would obtain, if the action were performed. There are cases, where Act-Consequentialism with this view on outcomes does not have the same implications as the Standard Version, if we accept certain views on welfare. (For example, some people think that a successful life that ends abruptly at a certain time is better than a life that is equally successful up to that time but that instead of ending abruptly at that time, fizzles out in mediocrity, even if the latter life contains more happiness.) However, this is not true of any of the cases I shall discuss here, even if we accept such a view of welfare. I could therefore – for the purposes of this thesis – just as well have chosen the view that we should only count the future of the possible world that would obtain, if the action were performed.³

There are various views as to what possible worlds are. The most plausible one is *Abstractionism*. This view claims that possible worlds are abstract things. More precisely, it says that possible worlds are states of affairs, i.e., a type of abstract entities which exist regardless of whether they obtain or not (much like propositions which exist regardless of whether they are true or

² See L. W. Sumner, *Welfare, Happiness, and Ethics*, Oxford University Press (1999).

³ For a discussion see, Erik Carlson, *Consequentialism Reconsidered*, Dordrecht, Kluwer Academic Publishers (1995), 48–58.

false). However, not every state of affairs is a possible world, according to Abstractionism. A possible world is rather a *maximally comprehensive* state of affairs. More precisely, it is a state of affairs such that, for any state of affairs, it either includes that state of affairs or its negation.

Now, there may seem to be some tension between the Standard Version and Abstractionism. As we shall see, the Standard Version claims (as almost everyone does) that actions are concrete occurrences, and does it not sound odd to say that the outcome of a concrete occurrence is something abstract? Well, it might sound odd if you assume that the relations between actions and outcomes are causal, but the Standard Version does not make that assumption. However, even if it made this assumption, this would not – as we shall see in Chapter Four when we discuss the Causal Consequence Version – be particularly embarrassing.

2.3 Actions

As I said, there are several views on what actions are and the Standard Version is not committed to any particular one. Still it is helpful to have a particular view in mind during our discussion. I shall therefore assume – following Alvin Goldman and others – that an action is (and that the Standard Version concurs that an action is) an instantiation of an action-property by an individual during an interval of time.⁴ The Standard Version is thus concerned with particular actions or act-tokens as contrasted with generic actions or act-types.

I do not have an account of when a property is an action-property. Of course, actions are intentional: A person does not act – does not instantiate an action-property – unless she does what she does intentionally. However, apart from that I do not have a lot to say about how to identify action-properties. Some examples of action-properties will have to do. “Running” is an obvious example of an action-property. “Reading” is another. Gertrude’s running between ten and ten thirty yesterday evening is thus an example of a particular action and Bertram’s reading between one and ten past one this afternoon is an example of another one.

The individuation of actions is a contentious matter. Suppose that Gertrude runs fast between ten and ten fifteen. Some would say that Gertrude’s *running fast* between ten and ten fifteen is a different action from Gertrude’s *running* between ten and ten fifteen. I do not have anything interesting to say about this issue. But I am not going to assume that simple actions are quite as fine-grained as that. Neither am I going to assume that actions are as modally fragile as this view might suggest: An action could have been performed in a slightly different way and at a slightly different time. (Although it is difficult to see how anyone could have a good reason to claim that actions are very

⁴ Alvin Goodman, *A Theory of Human Action*, Princeton University Press (1970).

fine-grained and at the same time deny that they are modally fragile, such view is of course logically possible.)

The characterisation of an action I gave above is a characterisation of a *simple* action. There are also combinations of simple actions. Some of these combinations are arguably actions in their own right, or as I shall say *compound* actions. Furthermore, some compound actions are *group* actions. A group action is a compound action where different agents perform (at least some of) the simple actions of the compound. Gertrude's saving a child from drowning between seven and ten past seven yesterday morning *and* Bertram's helping Gertrude saving a child from drowning between seven and ten past seven yesterday morning may well be an example of a group action.

I shall assume that the Standard Version is also concerned with compound actions. Moreover, I should perhaps stress that compound actions are particular actions too. Consequently, they should not be confused with action-types.

I said above that a person does not act unless her intentions are appropriately involved. It also seems true that a group does not act unless the intentions of the members are appropriately involved. I think it is clear that this at least requires that the members of the group are aware of one another. I am not alone here. Many believe that the intentions of the members are appropriately involved just in case the members jointly intend some aim; and that they do so just in case (a) each group member intends that they together pursue the aim, (b) each group member intends to do her part, (c) each group member forms these intentions – at least to some extent – because she believes that the other members also do so, and (d) each member believes that (a), (b), and (c) are met and that the other members believe this too.⁵

A few paragraphs back I said that some combinations of simple actions are actions themselves. Some people think that all combinations of simple actions compose a further action. The view I just embraced – i.e., that a combination of simple actions performed by different persons does not compose a group action unless the intentions of the persons performing the simple actions are appropriately involved – denies that all combinations of simple actions compose an action. According to the view I just embraced, there consequently are combinations of actions that are not compound actions. I shall think of combinations of actions that are not compound actions as *pluralities* of actions.

Finally, I shall for the sake of simplicity (awkwardly) count even a single simple action as a combination of actions.

⁵ See Michael E. Bratman, *Shared Agency: A Planning Theory of Acting Together*, Oxford University Press (2014). See also Christian List and Philip Pettit, *The Possibility, Design, and Status of Corporate Agents*, Oxford University Press (2011), 33–37.

2.4 Alternatives

The alternatives to a simple action are those actions that the agent could perform instead in the circumstances. I shall assume that we have a sufficiently good understanding of this notion. (In fact, it has proven to be surprisingly difficult to specify what the alternatives to an action are; there are several competing views. As I said, the Standard Version does not take side here, but this does not affect our business here.) Likewise, the alternatives to a combination of actions are those combinations the agent or agents could perform instead in the circumstances. Those are a function of what simple actions the agent or agents could perform instead in the circumstances.

For example, consider the combination “my shooting Victim, and your shooting Victim”. What are the alternatives to this combination? The answer to this question obviously depends on what I could do instead of shooting and on what you could do instead of shooting. To simplify things, suppose that I could only stand motionless and that you could only do so too. In that case the alternatives to the combination “my shooting Victim, and your shooting Victim” are the following three: “my shooting Victim, and your standing motionless”, “my standing motionless, and your shooting Victim”, and “my standing motionless, and your standing motionless.”

There would of course be a lot more alternatives to the combination in question if you and I had more alternatives. Suppose for example that I could also pick flowers. In that case, there would be two further alternatives, namely: “my picking flowers, and your shooting Victim”, and “my picking flowers, and your standing motionless.”

Let us now turn to combinations of three actions. Suppose that Gertrude also shoots Victim. Consider the combination of our three actions: “my shooting Victim, your shooting Victim, and Gertrude’s shooting Victim.” Suppose moreover (again in order to simplify things) that each of us could only stand motionless apart from shooting Victim. In that case, there are seven alternatives that may be ordered in three groups. First, none of us shoots Victim: i.e., “my standing motionless, your standing motionless, and Gertrude’s standing motionless” (this is one alternative). Second, two of us shoot Victim and one of us does not: i.e., “my shooting Victim, your shooting Victim, and Gertrude’s standing motionless”; “my standing motionless, your shooting Victim, and Gertrude’s shooting Victim”; “my shooting Victim, your standing motionless, and Gertrude’s shooting Victim” (that’s three alternatives). Third, one of us shoots Victim and two of us do not: i.e., “my shooting Victim, your standing motionless, and Gertrude’s standing motionless”; “my standing motionless, your shooting Victim, and Gertrude’s standing motionless”; “my standing motionless, your standing motionless, and Gertrude’s standing motionless” (that’s also three alternatives).

2.5 Negative Difference Makers

A distinctive feature of the Standard Version is that it implies that it is the difference an action makes to the total outcome that determines its moral status. An action that is wrong, for example, makes a difference for the worse. I am going to call actions that are wrong according to the Standard Version, “negative difference makers.” It is helpful to have a name of these actions since – as we shall see in section 3 – the cases we are concerned with here contain actions that are related in such a way that each action would have been a negative difference maker, if it were not for the other or others.

To get a better understanding of this feature of the Standard Version, it may be helpful to contrast the Standard Version with a view that Derek Parfit calls “the Share-of-the-Total View” (which he rejects).⁶ The Standard Version implies that it is the difference an action makes to the total outcome that determines its moral status. The Share-of-the-Total View, on the other hand, claims that it is the size of the share that an action contributes to the total outcome that determines its moral status. The difference between these two views is best brought out by way of an example.

Suppose that I could either rescue fifty persons together with you or rescue ten others by myself. Furthermore, suppose that you would rescue the fifty regardless of what I do. Finally, suppose that I join you rescuing the fifty. In that case, I act wrongly, according to the Standard Version but rightly according to the Share-of-the-Total View. I act wrongly according to the Standard Version because there is an alternative whose outcome is intrinsically better. It would have been better if I had rescued the ten by myself since more persons would then have been saved (sixty instead of fifty). I act rightly according to the Share-of-the-Total View since my total share of people saved is larger (it is twenty-five of fifty) than it would be if I had gone off and saved the ten by myself (it would then be ten of sixty). In this case, I think it is pretty evident that the Standard Version does better than the Share-of-the-Total View: Ten more people would have been saved, if I instead had acted in accordance with the Standard Version.

3. Redundant Negative Difference Making

In this section I shall stipulate under what circumstances an action is what I shall call a “redundant negative difference maker.” I shall go on to show that our actions in Case One and Case Two are redundant negative difference makers. This is important – as I indicated in the introduction to this chapter –

⁶ Parfit compares the Share-of-the-Total View with a view similar to the Standard Version in order to show that the latter is preferable to the former. He also considers a revised version of the Share-of-the-Total View and rejects that too. Derek Parfit, *Reasons and Persons*, Oxford, Clarendon Press (1984), 67–69.

because a proponent of Act-Consequentialism might, in light of these cases, abandon the Standard Version and instead suggest that an action is wrong in virtue of having an alternative whose outcome is intrinsically better *or* being a redundant negative difference maker. As I mentioned, I shall defend a proposal along those lines in Chapter Five and Chapter Six.

I shall proceed as follows. In section 3.1 I say under what circumstances I take actions to be redundant negative difference makers. I go on to show that our actions in Case One and Case Two are redundant negative difference makers. In Section 3.2, I suggest that two other cases that are very similar to Case One and Case Two but differ in that there are more people involved also challenge the Standard Version. I go on to show that the actions involved in these cases are redundant negative difference makers as well. Finally, in section 3.3 I suggest that a case that is similar to Case Two but differ in that that my action depends counterfactually on your action also challenge the Standard Version. I go on to show that our actions are redundant negative difference makers here too. I also address the issue that my characterisation of redundant negative difference making implies that some actions that do not seem to be redundant still come out as redundant negative difference makers.

3.1 Redundant Negative Difference Makers

So, under what circumstances is an action a redundant negative difference maker? Let us begin by considering Case One again:

Case One

You and I shoot and kill Victim. Our bullets arrive at the same time and each shot would have killed Victim by itself. Victim would not have been killed, if neither of us had pulled the trigger. Apart from shooting Victim, you and I have only one alternative each: each could also have abstained from shooting Victim. Furthermore, each of us would have performed his or her alternative – i.e., each of us would have abstained from shooting Victim – if he or she had not shot him. We act independently of one another, Victim’s level of well-being is not affected by either of the shots taken by itself (e.g., Victim does not experience any pain), and no one else is affected. Finally, Victim would have gone on to live a life well worth living, if neither you nor I had shot him.

What is true of our actions here? There are particularly three things that come to mind. First, neither of us has an alternative such that the outcome would be intrinsically better if that alternative and the other person’s action had been performed. Second, each of us has an alternative such that the outcome would not be intrinsically worse if that alternative and the other person’s action had

been performed. Third, there is an alternative to the combination of your action and my action such that the outcome would be intrinsically better if that alternative had been performed.

I shall call actions of which this is true “redundant negative difference makers.”⁷ Why is that a fitting name? Both actions are such that they would have been a negative difference maker, if it were not for the other. For example, in the closest possible world where I do not shoot Victim and you shoot him, you have an alternative whose outcome would be intrinsically better, namely, abstaining from shooting him.

With these remarks in mind, I stipulate this:

Redundant Negative Difference Making

An action *a* is a redundant negative difference maker, if and only if, there is a combination of actions *C* such that

- (1) there is no alternative to *a* such that the outcome would be intrinsically better if that alternative and the combination *C* had been performed, and there is at least one alternative to *a* such that the outcome would not be intrinsically worse if that alternative and the combination *C* had been performed,
- (2) there is no alternative to the combination *C* such that the outcome would be intrinsically better if that alternative and *a* had been performed, and there is at least one alternative to the combination *C* such that the outcome would not be intrinsically worse if that alternative and *a* had been performed,
- (3) there is an alternative to the combination *a* and *C* such that the outcome would be intrinsically better if that alternative had been performed.

In Case One there are two simple actions to consider. However, there are obviously also (as we shall see in the next section) the same type of cases with more than two simple actions to consider. Hence, I wrote “there is a combination of actions *C* such that” instead of “there is an action *b* such that.” As I said, I also refer to a simple action as a combination of actions. So, the variable *C* may also pick out a simple action.

Are both my action and your action, respectively, redundant negative difference makers in Case One, according to Redundant Negative Difference Making? Yes, consider my action, for instance. There is a combination of actions – namely, your action – such that the three clauses of Redundant Negative Difference Making are satisfied. Clause (1) is satisfied since I only have one alternative – i.e., abstaining from shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse if I had performed that alternative and you had shot Victim since one shot is sufficient

⁷ Christian List suggested that I call these things “redundant negative difference makers.”

to kill him. Moreover, clause (2) is satisfied since you only have one alternative – i.e., abstaining from shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse if you had performed that alternative and I had shot him since (again) one shot is sufficient to kill him. Finally, clause (3) is satisfied since we have an alternative – i.e., the combination “my abstaining from shooting Victim, and your abstaining from shooting Victim” – such that the outcome would have been intrinsically better, if we had performed that combination: Victim would under those circumstances not have been killed.

A parallel argument establishes that your action is a redundant negative difference maker here, according to Redundant Negative Difference Making.

Consider now Case Two:

Case Two

You and I shoot one shot each at Victim. My bullet arrives first and kills Victim. Your bullet arrives very shortly after mine but does not kill Victim since my bullet has already done that. Your bullet would however have killed Victim, if my bullet had not done so first. Victim would not have been killed, if neither of us had pulled the trigger. Apart from shooting Victim, you and I have only one alternative each: each could also have abstained from shooting Victim. Furthermore, each of us would have performed his or her alternative – i.e., each of us would have abstained from shooting Victim – if he or she had not shot him. We act independently of one another, Victim’s level of well-being is not affected by my shot (e.g., Victim does not experience any pain), and no one else is affected. Finally, Victim would have gone on to live a life well worth living, if neither you nor I had shot him, and the extra time Victim would have lived, if I had not shot him would have been neither intrinsically good nor intrinsically bad for him.

Both our actions are redundant negative difference makers in Case Two also, according to Redundant Negative Difference Making. Consider your action this time. There is a combination – namely, my action – such that the clauses of Redundant Negative Difference Making are satisfied. Clause (1) is satisfied since you only have one alternative – i.e., abstaining from shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse if you had abstained from shooting Victim and I had shot him since one shot is sufficient to kill him. Furthermore, clause (2) is satisfied since I only have one alternative – i.e., abstaining from shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse if I had abstained from shooting Victim and you had shot Victim since (once more) one shot is sufficient to kill him. Finally, clause (3) is satisfied since we have an alternative – i.e., the combination “my abstaining from shooting Victim,

and your abstaining from shooting Victim” – such that the outcome would have been intrinsically better, if we had performed that combination: Victim would under those circumstances not have been killed.

A parallel argument establishes that my action is a redundant negative difference maker in Case Two, according to Redundant Negative Difference Making.

3.2 Cases Involving More Than Two Actions

However, there are also cases where there are more than two actions involved.

The following case is such a case:

Case Three

You, Gertrude, and I shoot and kill Victim. Our bullets arrive at the same time and each shot would have killed Victim by itself. Victim would not have been killed, if none of us had pulled the trigger. Apart from shooting Victim, each of us has only one alternative: each could also have abstained from shooting Victim. Furthermore, each of us would have performed his or her alternative – i.e., each of us would have abstained from shooting Victim – if he or she had not shot him. We act independently of one another, Victim’s level of well-being is not affected by either of the shots taken by itself (e.g., Victim does not experience any pain), and no one else is affected. Finally, Victim would have gone on to live a life well worth living, if none of us had shot him.

According to the Standard Version, none of us acts wrongly in Case Three, since none of us has an alternative whose outcome is intrinsically better. Intuitively, each of us does act wrongly, however. (If you agree that each of us do act wrongly in Case One, it is difficult to see how you could deny that each of us acts wrongly here.) So, a proponent of Act-Consequentialism that want to suggest that an action might also be wrong in virtue of being a redundant negative difference maker would presumably want to show that our actions in Case Three are redundant negative difference makers.

Are our actions in Case Three redundant negative difference makers, according to Redundant Negative Difference Making? Yes, they are. Consider my action once more. There is a combination of actions – namely, your action and Gertrude’s action – such that the relevant clauses are satisfied. Clause (1) is satisfied since I only have one alternative – i.e., abstaining from shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse if I had performed that alternative and you and Gertrude had shot Victim since one shot is sufficient to kill him. Clause (2) is also satisfied since although you and Gertrude together have several alternatives, the outcome would have been neither intrinsically better nor intrinsically worse if you and Gertrude had performed any of them and I had shot Victim since

(again) one shot is sufficient to kill him. Finally, clause (3) is satisfied since you, Gertrude and I have an alternative – i.e., the combination “my abstaining from shooting Victim, your abstaining from shooting Victim, Gertrude’s abstaining from shooting Victim” – such that the outcome would have been intrinsically better, if we had performed that combination: Victim would under those circumstances not have been killed.

Parallel arguments show that your action and Gertrude’s action also are redundant negative difference makers in this case, according to Redundant Negative Difference Making.

Consider now this case:

Case Four

You, Gertrude and I together shoot and kill Victim. Our bullets arrive at the same time but we are dealing with a very odd victim: he won’t die unless two bullets hit him simultaneously (but two bullets would do the job just fine). Victim would not have been killed, if none of us had pulled the trigger. Apart from shooting Victim, each of us has only one alternative: each could also have abstained from shooting Victim. Furthermore, each of us would have performed his or her alternative – i.e., each of us would have abstained from shooting Victim – if he or she had not shot him. We act independently of one another, Victim’s level of well-being is not affected by either of the shots taken by itself (e.g., Victim does not experience any pain), and no one else is affected. Finally, Victim would have gone on to live a life well worth living, if none of us had shot him.

According to the Standard Version, none of us acts wrongly in this case either, since, again, none of us has an alternative whose outcome is intrinsically better. However, it seems intuitively clear in this case too, that we do act wrongly. Consequently, a proponent of Act-Consequentialism that want to suggest that an action also is wrong in virtue of being a redundant negative difference maker would likely want to show that our actions are redundant negative difference makers here too.

Our actions are redundant negative difference makers in Case Four, according to Redundant Negative Difference Making. Focus on my action once again. When I discussed Case Three a few paragraphs back, I took the variable *C* to pick out both your action and Gertrude’s action. Should I do so in this case too? No! Clause (2) would not then be satisfied. You and Gertrude have an alternative – namely, “your abstaining from shooting Victim, and Gertrude’s abstaining from shooting Victim” – such that the outcome would be intrinsically better if you and Gertrude had performed that alternative and I had shot Victim since it takes at least two shots to kill him in Case Four.

Let the variable *C* instead pick out, for instance, Gertrude’s action. In that case, the clauses of Redundant Negative Difference Making are satisfied. Clause (1) is satisfied since I only have one alternative – i.e., abstaining from

shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse if I had done that and Gertrude had shot Victim since he still gets killed under those circumstances: Your shot does not vanish (it is rather part of the circumstances), and Gertrude’s shot together with your shot is sufficient to kill Victim. Moreover, clause (2) is satisfied since Gertrude only has one alternative – i.e., abstaining from shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse if she had done so and I had shot him since Victim still gets killed in those circumstances: Again, your shot does not just disappear. Finally, clause (3) is satisfied since Gertrude and I have an alternative – i.e., the combination “Gertrude’s abstaining from shooting Victim, and my abstaining from shooting Victim” such that the outcome would have been intrinsically better, if we had performed that combination: Victim would under those circumstances not have been killed since it takes at least two shots to kill him.

Once more, parallel arguments demonstrate that your action and Gertrude’s action also are redundant negative difference makers in Case Four, according to Redundant Negative Difference Making.

3.3 A Case Involving Dependence

There are also cases like Case One and Case Two, but where one of the actions depend on the other.

Consider this case, for example:

Case Five

This is exactly like Case Two except that my shot counterfactually depends on your shot. (I am simply the sort of person that wouldn’t shoot someone unless someone else goes first.) We may imagine that you did shoot first, but my bullet nevertheless reached Victim before yours.

Case Five also challenges the Standard Version. My action is not wrong here, according to the Standard Version, since I do not have an alternative whose outcome is intrinsically better: Victim would still have been killed by you, if I had not shot him. Intuitively, I do act wrongly, however. If I act wrongly in Case Two where my action does not depend on yours, it is difficult to see why I would not be acting wrongly in Case Five where my action does depend on yours (perhaps particularly so when my character seems shady). So, again, a proponent of Act-Consequentialism that want to suggest that an action also is wrong in virtue of being a redundant negative difference maker would presumably want to show that my action is a redundant negative difference maker in Case Five as well.

Well, my action is a redundant negative difference maker in Case Five, according to Redundant Negative Difference Making. There is a combination

– your action – such that the relevant clauses are satisfied. Clause (1) is satisfied since I only have one alternative – i.e., abstaining from shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse if I had done that and you had shot Victim since your shot is sufficient to kill him. Clause (2) is satisfied too since you only have one alternative – i.e., abstaining from shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse if you had done so and I had shot him since (again) one shot is sufficient to kill him. (This is true although I would not have shot him unless you went first.) Finally, clause (3) is satisfied since we have an alternative – i.e., the combination “my abstaining from shooting Victim, and your abstaining from shooting Victim” such that the outcome would have been intrinsically better, if we had performed that combination: Victim would under those circumstances not have been killed.

A parallel argument demonstrates that your action is a redundant negative difference maker in Case Five too, according to Redundant Negative Difference Making. But this is not very attractive: Your action is not exactly redundant in the ordinary sense of “redundant” since Victim would not have been killed if you had not shot him. To avoid this breach of ordinary usage, I could add the following clause to the stipulation above: “(4) the outcome would have been neither intrinsically better nor intrinsically worse, if *a* had not been performed.” This clause would exclude your action since the outcome would have been better, if you had not shot Victim. However, nothing substantial hangs on this. So, to keep things simpler, I am not going to add this clause.

4. Redundant Causation

In this section I shall say under what circumstances an action is a redundant cause of another event. Moreover, I shall go on to show that the actions in Case One and Case Two are redundant causes. This is of interest to us – as I suggested in the introduction to this chapter – because a proponent of Act-Consequentialism might, in view of the challenge posed by Case One and Case Two, abandon the Standard Version and instead adopt the Causal Consequence Version of Act-Consequentialism. This version might seem promising since it is intuitively clear, I think, that you and I, respectively, cause Victim’s death in Case One, for example. A proponent of the Causal Consequence Version might argue that the fact that our actions are redundant causes of Victim’s death supports this intuition. She might want to suggest that an action is a cause, if among other things it is a redundant cause, for instance. As I said, I shall discuss this view and the Causal Consequences Version (which I reject) in Chapter Four.

I shall proceed as follows: In section 4.1 I say under what circumstances events are redundant causes. I then show that the actions in Case One and Case Two are redundant causes. In Section 4.2 I show that Case Three and Case

Four (i.e., cases that differ from Case One and Case Two in that they involve more people) also are cases of redundant causation. Recall that these cases pose the same kind of challenge as Case One and Case Two. In Section 4.3 I discuss Case Five (i.e. the case where my action depends on yours). Finally, in section 4.4, I distinguish various cases of redundant causation. They fall into two large groups: cases of over-determination (where Case One belongs) and cases of pre-emption (where Case Two belongs).

4.1 Redundant Causes

Cases of redundant causation may be characterised in several ways. For example, David Lewis characterises them as follows:

Suppose we have two events c_1 and c_2 , and another event e distinct from both of them; and in actuality all three occur; and if either one of c_1 or c_2 had occurred without the other, then also e would have occurred; but if neither c_1 nor c_2 had occurred then e would not have occurred. Then I shall say that c_1 and c_2 are *redundant causes* of e .⁸

It is clear that Case One, for example, is a case of redundant causation (i.e., that your action and my action are redundant causes of Victim's death), according to Lewis's account. My shot and your shot occur. Furthermore, Victim's death – an event distinct from our shots – occurs, and Victim's death would have occurred if my shot had occurred without your shot, and the other way around. Finally, Victim's death would not have occurred, if neither your shot nor my shot had occurred.

I am nevertheless going to characterise cases of redundant causation a bit differently. I have three reasons for that. First, Lewis only provides sufficient conditions, but I want to provide necessary conditions as well. Second, Lewis is concerned with cases involving two redundant causes, but I am concerned with cases that involves more than two redundant causes, e.g., Case Three and Case Four. Third, Lewis tells us under what circumstances several (or rather two) events are redundant causes, but I want to say under what circumstances a single event is a redundant cause. I want – for the sake of clarity – it to be similar to my stipulation of redundant negative difference makers.

I shall therefore make the following suggestion:

Redundant Causation

An event c is a *redundant cause* of an event e , if and only if, there are events $c_1 \dots c_n$ such that

⁸ David Lewis, "Postscripts to "Causation"," 172–213 in David Lewis *Philosophical Papers Vol. II*, Oxford, Oxford University press 1986, 193.

- (1) e would have occurred, if c had not occurred and $c_1 \dots c_n$ had occurred
- (2) e would have occurred, if none of $c_1 \dots c_n$ had occurred and c had occurred
- (3) e would not have occurred, if neither c nor none of $c_1 \dots c_n$ had occurred.⁹

Let us now consider Case One (i.e., the case where you and I shoot Victim and our bullets arrive at the same time). Are my action and your action, respectively, redundant causes of Victim's death, according to the suggestion I made, i.e., Redundant Causation? Yes, consider your action, for example. Clause (1) is satisfied because Victim's death would have occurred, if you had not shot Victim and I had shot him since one shot is enough to kill him. (The variables $c_1 \dots c_n$ may pick out only one event.) Similarly, clause (2) is satisfied because Victim's death would have occurred, if I had not shot Victim and you had done so since one shot is sufficient to kill him. Finally, clause (3) is satisfied since Victim's death would not have occurred, if neither you nor I had shot Victim since it takes at least one of our shots to kill him. Consequently, your action is a redundant cause of Victim's death, according to Redundant Causation.

A parallel argument establishes that your action is a redundant cause of Victim's death, according to Redundant Causation. Moreover, an analogous argument establishes that your action and my action, respectively, are redundant causes of Victim's death in Case Two (i.e., the case where you and I shoot Victim but my bullet arrives before yours), according to Redundant Causation.

There are a few things I would like to discuss before I move on. First, the vocabulary used here (inherited from Lewis and now established) is slightly odd. It is for example odd to call your action in Case Two a "redundant cause" of Victim's death since it is not a cause at all. I shall nevertheless go on to use this vocabulary. It is established and I cannot find a better way to talk about these things.

Second, it might be argued that, for example, your action in Case One is not a redundant cause of Victim's death. Victim's death in Case One would plausibly have been slightly different, if you had not shot him since one bullet (yours) would then have pierced his body instead of two, and this might be enough to count that (counterfactual) death as a numerically different death. Moreover, it may also be argued that my action in Case Two is not a redundant cause of Victim's death either. Victim's death in Case Two would have happened slightly later, if I had not shot him since your bullet would have arrived a bit later, and this might also be sufficient to count this death as a numerically different death. However, just as I assumed (in section 2.3) that an action could have been performed in a slightly different way or at a slightly different time

⁹ I say "nor none of $c_1 \dots c_n$ had occurred" rather than, "nor $c_1 \dots c_n$ had occurred" because the closest world where $c_1 \dots c_n$ do not occur may be a world where some of $c_1 \dots c_n$ occur. I want to exclude that.

(or both) I am going to assume that an event could have happened in a slightly different way or at a slightly different time (or both). After all, actions are events.

Third, note that even if it were the case that neither Case One nor Case Two is a case of redundant causation they would still be cases of redundant negative difference making. This is determined by the value assumptions involved: e.g., that only Victim is affected, that Victim would have gone on to live a life well worth living, if he had not been attacked by us, and (with respect to Case Two) that the extra time Victim would have lived in if I had not shot him would be neither good nor bad for him.

4.2 Case Three and Case Four

Let's start with Case Three (i.e., the case where there are three shooters – you, me and Gertrude – and where each shot is sufficient to kill Victim.) This is a case of redundant causation. My action, your action, and Gertrude's action, respectively, are redundant causes of Victim's death, according to Redundant Causation.

Consider my action, for instance. Clause (1) is satisfied because Victim's death would have occurred, if I had not shot Victim and you and Gertrude had shot him since one shot is enough to kill him. (Here $c_1 \dots c_n$ pick out your action and Gertrude's action). Moreover, clause (2) is satisfied because Victim's death would have occurred, if neither you nor Gertrude had shot Victim and I had done so since (again) one shot is sufficient to kill him. Finally, clause (3) is satisfied since Victim's death would not have occurred, if neither you nor I nor Gertrude had shot Victim since it takes at least one shot to kill him. Consequently, my action is a redundant cause of Victim's death.

Parallel arguments show that your action and Gertrude's action, respectively, also are redundant causes of Victim's death in Case Four, according to Redundant Causation.

Consider now Case Four (i.e., the case where there are three shooters – you, me and Gertrude – but where two shots are needed to kill Victim.) This is also a case of redundant causation: My action, your action, and Gertrude's action, respectively, are redundant causes of Victim's death here too, according to Redundant Causation.

Consider my action again. When I discussed Case Three, I took the variables $c_1 \dots c_n$ to pick out both your action and Gertrude's action. Should I do so in this case too? No, clause (2) would not be satisfied under these circumstances. It is not the case that Victim's death would have occurred if neither you nor Gertrude had shot Victim, and I had done so, since one shot is not enough to kill Victim here.

However, let the variables $c_1 \dots c_n$ instead pick out, for example, Gertrude's action. Clause (1) is then satisfied: Victim's death would have occurred, if I had not shot Victim and Gertrude had shot him since Gertrude's shot together

with your shot is sufficient to kill Victim. Your shot does not just disappear; it is rather part of the circumstances. Clause (2) is also satisfied since Victim's death would still have occurred, if Gertrude had not shot Victim and I had shot him since my shot together with your shot is sufficient to kill Victim. Again, your shot does not just disappear. Finally, clause (3) is satisfied since Victim would not have been killed, if neither I nor Gertrude had shot Victim: Two shots are necessary to kill him. So, my action is a redundant cause of Victim's death here too, according to Redundant Causation.

Parallel arguments show that your action and Gertrude's action, respectively, also are redundant causes of Victim's death in Case Four, according to Redundant Causation.

4.3 Case Five

Your action and my action are redundant causes of Victim's death in Case Five (i.e., the case where my action depends on yours), according to Redundant Causation.

Consider my action once again. Clause (1) is satisfied since Victim's death would have occurred, if my shot had not occurred and your shot had occurred. Furthermore, clause (2) is satisfied since Victim's death would also have occurred, if your shot had not occurred and my shot had occurred. (The fact that I would not have shot Victim unless you had done so does not imply that it is false that Victim's death would have occurred, if your shot had not occurred and my shot had occurred.) Finally, clause (3) is satisfied since Victim's death would not have occurred if neither your shot nor my shot had occurred. My action is thus a redundant cause of Victim's death in Case Five as well, according to Redundant Causation.

A parallel argument shows that your action also is a redundant cause of Victim's death in Case Five, according to Redundant Causation. However, just as it is a breach of the ordinary sense of "redundant" to say that your action is a "redundant difference maker", it is a breach of the ordinary sense of "redundant" to say that your action is a "redundant cause". Again, Victim would not have been killed if you had not shot him. To avoid this breach of ordinary usage, I could add a clause to Redundant Causation. I could add the following one: "(4) *e* would have occurred, if *c* had not occurred". This would exclude your action since Victim's death would not have occurred, if you had not shot Victim. But nothing substantial turns on this here either. I shall not therefore add this clause.

I noted above that my characterisation of redundant difference making has the upshot that many actions that are not redundant in the ordinary sense of "redundant" come out as redundant negative difference makers. Likewise, my characterisation of redundant causation has the upshot that many actions that are not redundant in the ordinary sense of "redundant" come out as redundant causes.

For example, suppose that I fall off a very high building and collide with the ground. Furthermore, suppose that the collision kills me. In that case it may well be the fall and the collision, respectively, are redundant causes of my death. Consider the fall. Clause (1) seems to be satisfied since my death would presumably have occurred, if the fall had not occurred and the collision had occurred. (Another very similar fall but numerically distinct would probably have been responsible for the collision then.) Furthermore, clause (2) is arguably satisfied since my death would likely have occurred if the fall had occurred but not the collision. (I am falling off a high building.) Finally, clause (3) seems to be satisfied since my death would not have occurred, if neither the fall nor the collision would have happened. So, the fall is a redundant cause of my death here. But that does not seem to be quite right. (A parallel argument shows that the collision is a redundant cause of my death, according to Redundant Causation.) But as I said, nothing substantial is at stake here, so I shall leave things as they are.

4.4 Over-Determination and Pre-Emption

What are the differences between Case One (i.e., the case where you and I shoot Victim and our bullets arrive at the same time) and Case Two (i.e., the case where you and I shoot Victim but my bullet arrive before yours)? One difference is that in Case One either both our actions are causes or neither of them is since there is no relevant difference between them; whereas in Case Two it is clear that only my action is a cause.

A further difference is that the causal processes running from our actions in Case One are not interrupted; whereas the causal process running from your action is interrupted in Case Two (but not the one running from mine). The causal process running from your action is interrupted by Victim's death.

Accordingly, I shall say that a case of redundant causation is a case of over-determination if and only if either both redundant causes are causes or neither of them is a cause, and the causal process running from each redundant cause proceeds without being interrupted. Furthermore, I shall say that a case of redundant causation is a case of pre-emption if and only if the causal process running from the pre-empting redundant cause is not interrupted and the causal process running from the pre-empted one is interrupted. (I am for the sake of simplicity considering cases involving two redundant causes.)

Consider now this case:

Case Six

You and I throw a brick each at a window. My brick shatters the window. However, your brick would have shattered the window, if I had not thrown my brick. Immediately after our bricks were thrown, my brick (it's an odd

brick) transmitted a signal to your brick (which it wouldn't have done, if I had not thrown my brick) that sent your brick off course.

This case is a case of redundant causation, according to Redundant Causation: the shattering would have occurred if my throw had occurred without yours, and the other way around; but it would not have occurred if neither of our throws had occurred. It also seems clear that my throw is a cause and that your throw is not. Moreover, the causal process running from my throw is not interrupted; whereas the causal process running from your throw is. So, Case Six seems to be a case of pre-emption. However, there is a striking difference between Case Six and Case Two. In Case Six the causal process running from your action is interrupted before the effect (i.e., the shattering of the window) occurs; not as in Case Two where the causal process running from your action is interrupted when the effect (i.e. Victim's death) occurs.

In light of these remarks, I shall say that a case of pre-emption is a case of *late* pre-emption if and only if the causal process running from the pre-empted alternative is interrupted when the effect occurs. Moreover, I shall say that a case of pre-emption is a case of *early* pre-emption if and only if the causal process running from the pre-empted alternative is interrupted before the effect occurs.¹⁰

¹⁰ I am following Lewis in this section. See David Lewis "Causation as Influence," 74–106 in John Collins, Ned Hall, L. A. Paul (ed.), *Causation and Counterfactuals*, Cambridge Massachusetts, MIT Press 2004, 80. See also David Lewis, "Postscripts to "Causation"," 172–213.

Chapter Three, Replies

1. Introduction

In the previous chapter, I carefully described the Standard Version of Act-Consequentialism and Case One and Case Two, i.e., the cases where you and I shoot and kill a third person called “Victim”. I noted that they both are cases of redundant negative difference making and that Case One is a case of causal over-determination and that Case Two is a case of pre-emption. As I noted in the chapter before that – i.e., Chapter One – both Case One and Case Two challenge the Standard Version. It implies that neither you nor I act wrongly in these two cases. Intuitively, it seems that each of us does, however.

In this chapter, I shall consider some replies that the proponent of the Standard Version might make.¹¹ There are two broad strategies that she might employ. First, she may concede that the implications of the Standard Version in Case One and in Case Two – i.e., that neither you nor I act wrongly – are counterintuitive, and try to find a way to deal with that. For example, she might try to explain away this intuition. Second, she may deny that the implications of the Standard Version in Case One and in Case Two are counterintuitive. She could claim that we have the less specific intuition that something is not as it should be in these cases rather than the intuition that you and I, respectively, act wrongly. She could then come up with a suggestion of what it is that is not as it should be in these cases. I shall suggest that although both these strategies have something going for them, they are sufficiently problematic to motivate an investigation into whether there are any plausible versions of Act-Consequentialism that have the intuitively correct implications in Case One and Case Two. To be clear, I am not suggesting that these two strategies fail. I am merely suggesting that they are sufficiently problematic to motivate further investigation.

I shall proceed as follows: In section 2, I briefly return to the challenge that Case One and Case Two pose to the Standard Version. One issue is here whether it is plausible to say that both of us act wrongly in Case Two (i.e., the pre-emption case). In section 3, I argue that the first strategy that the proponent

¹¹ As I said in the introduction to this thesis, a consequentialist might abandon the Standard Version and instead adopt Rule-Consequentialism or Consequentialist Generalisation to deal with Case One and Case Two. But as I also said, I am concerned with what a proponent of the Standard Version might say as regards these cases in this thesis.

of the Standard Version might employ is problematic, i.e., the one that involves conceding that the implications of the Standard Version in Case One and in Case Two are counterintuitive. In section 4, I briefly introduce the second strategy, i.e., the strategy that involves denying that the implications of the Standard Version in Case One and in Case Two are counterintuitive and tries to find something else that is not as it should be in these cases. There are primarily three suggestions of what that is in the literature: that you and I together act wrongly, that you and I respectively act subjectively wrongly, and that you and I are uncooperative. In sections 5–7, I consider these candidates in the order just mentioned.

2. The Challenge

Let us refresh our memories of the challenge that Case One and Case Two pose to the Standard Version. I also want to note an important difference between Case One and Case Two. Here is Case One again:

Case One

You and I shoot and kill Victim. Our bullets arrive at the same time and each shot would have killed Victim by itself. Victim would not have been killed, if neither of us had pulled the trigger. Apart from shooting Victim, you and I have only one alternative each: each could also have abstained from shooting Victim. Furthermore, each of us would have performed his or her alternative – i.e., each of us would have abstained from shooting Victim – if he or she had not shot him. We act independently of one another, Victim’s level of well-being is not affected by either of the shots taken by itself (e.g., Victim does not experience any pain), and no one else is affected. Finally, Victim would have gone on to live a life well worth living, if neither you nor I had shot him.

Neither you nor I act wrongly in Case One, according to the Standard Version. There is no alternative to what I do whose outcome is intrinsically better: If I had not shot Victim you would still have done so. This is true of what you do as well. However, it seems intuitively clear that you and I, respectively, act wrongly in Case One. Victim is deprived of a future that would have been well worth living as a consequence of what we do. Furthermore, each of our actions is causally sufficient for killing Victim, and each of us would also have killed him, even if the other had abstained from shooting.

Consider now Case Two:

Case Two

You and I shoot one shot each at Victim. My bullet arrives first and kills Victim. Your bullet arrives very shortly after mine but does not kill Victim since my bullet has already done that. Your bullet would however have killed Victim, if my bullet had not done so first. Victim would not have been killed, if neither of us had pulled the trigger. Apart from shooting Victim, you and I have only one alternative each: each could also have abstained from shooting Victim. Furthermore, each of us would have performed his or her alternative – i.e., each of us would have abstained from shooting Victim – if he or she had not shot him. We act independently of one another, Victim's level of well-being is not affected by my shot (e.g., Victim does not experience any pain), and no one else is affected. Finally, Victim would have gone on to live a life well worth living, if neither you nor I had shot him, and the extra time Victim would have lived, if I had not shot him would have been neither intrinsically good nor intrinsically bad for him.

Neither of us acts wrongly in Case Two either, according to the Standard Version. It is true in this case too that neither of us has an alternative whose outcome is intrinsically better since Victim would have been killed if I had not shot him, and since he would still have been if you had not shot him.

I take it that it is intuitively clear that I act wrongly in Case Two. However, is it really plausible to say that you act wrongly here? This case is a case of pre-emption: Your bullet reaches Victim when he is already dead. Well, I think that a moral theory at least may imply that you act wrongly in Case Two. Note that it is true also in this case that Victim is deprived of a future that would have been well worth living as a consequence of what we do, that each of our actions is causally sufficient for killing him, and that each of us would have killed him, if the other had abstained from shooting. On account of this, it is not clear (I think) that there is a morally relevant difference between what you and I do in Case Two. So, it is not clear (I think) that it is morally relevant that it is I and not you that kill Victim in Case Two, for example. In any event, so I shall assume. (I shall briefly return to this issue in Chapter Six, section 2.2.)

3. Conceding the Intuition

In the introduction of this chapter I said that there are two broad strategies that the proponent of the Standard Version may employ. In this section I discuss the first of these two strategies, i.e., the one that involves conceding that the

implications of the Standard Version in Case One and in Case Two are counterintuitive. If the proponent of the Standard Version concedes that we have the intuition that you and I, respectively, act wrongly in Case One there seem to be two options open to her. First, she could deny that this intuition has any evidential value. Second, she could claim that other serious contenders – she might even be prepared to say all of them – either are partners in crime (i.e., they all stumble on Case One and Case Two) or have other just as serious difficulties to deal with. I am going to argue that both these replies are unsatisfactory.

I discuss the first option in section 3.1, and the second, in section 3.2.

3.1 Denying the Evidential Value of the Intuition

The proponent of the Standard Version who wishes to deny the evidential value of the intuition that you and I act wrongly in Case One and Case Two seems to have two options. First, she may deny that moral intuitions – generally speaking, and including the intuition that you and I respectively act wrongly in Case One and Case Two – have evidential value. Second, she might deny that the *particular* intuition that you and I respectively act wrongly in Case One and Case Two has evidential value.¹²

There are several reasons why the first option is problematic. For one thing, it is difficult to see how normative ethics could get by without moral intuitions. Moreover, the view that the Standard Version is a better theory than competing moral theories arguably rests in part on moral intuitions. How could we otherwise rank it above competing moral theories that are equally theoretically virtuous? (For instance, how could we rank it above a theory that claims that an action *a* is *right*, if and only if, there is an alternative to *a* whose consequences are intrinsically better than those of *a*, that an action *a* is *obligatory*, if and only if, the outcome of *a* is intrinsically worse than the outcome of any of its alternatives would be, and that an action *a* is *wrong*, if and only if, there is no alternative to *a* whose outcome would be intrinsically better than that of *a*?)

The second option, i.e., denying that the *particular* intuition that you and I respectively act wrongly in Case One and Case Two has evidential value, is also problematic. It seems plausible that the proponent of the Standard Version in order to deny the evidential value of this intuition must deny that the

¹² Frank Jackson has suggested that a proponent of the Standard Version should deny that you and I act wrongly in Case One and Case Two because “it runs counter to the whole trust of consequentialist thinking about morality.” (He has act-consequentialism in mind.) However, as I shall argue in Chapter Five, section 2.5, I do not think that this is true. See Frank Jackson, “Which Effects?”, *Reading Parfit*, Blackwell (1997), 47–48.

best explanation of people's having the intuition is that you and I respectively act wrongly in these cases. Arguably, an intuition has evidential value just in case the best explanation of why people have the intuition is that it is true. However, it seems very difficult to do so without also denying that intuitions supporting the Standard Version have evidential value.

For example, suppose that the proponent of the Standard Version tries to explain away the intuition that you and I respectively act wrongly in Case One and Case Two by appealing to natural selection. I have no idea whether there is a good evolutionary explanation for that intuition, but let us assume that the proponent of the Standard Version has one. Suppose she claims that individuals that have qualms about inflicting violence on others are more successful in spreading their genes. However, the proponent of the Standard Version arguably has to explain why intuitions that support the Standard Version do not meet a similar fate. Generally, if the theory of evolution is correct and explains why some intuitions are not truth tracking, we arguably need an explanation why others are.

This is of course inconclusive. We may find, if we study things a bit closer, that evolution best explains the intuition that you and I respectively act wrongly in Case One but not the intuitions that support the Standard Version. Or we may find that our religious heritage best explains the intuition that you and I respectively act wrongly in Case One but not the intuitions that support the Standard Version. Or there may be another explanation altogether. However, it is fair to say that recent attempts at explaining why intuitions that do not conform to the Standard Version could be explained away without at the same time implicating intuitions that support the Standard Version have not been very persuasive.¹³ I therefore think that my skepticism here is justified.

3.2. Partners in Crime

As I said above, the proponent of the Standard Version may also suggest that the competing normative theories worth considering either are partners in crime or have other just as serious difficulties to deal with.

For example, take a view that says that you act wrongly only if you harm someone, and that you harm someone only if the person would have been intrinsically better off if the action had not been performed. This view implies that neither you nor I act wrongly in Case One and Case Two. Consider Case One, for instance. Victim would not have been intrinsically better off, if I had

¹³ Peter Singer, for example, has argued that Standard Version intuitions are different from other types of intuitions. Both Folke Tersman and Selim Berker have convincingly challenged this view. See Peter Singer, "Ethics and Intuitions," *Journal of Ethics* 9: 331–352 (2005); Selim Berker, "The Normative Insignificance of Neuroscience," *Philosophy & Public Affairs* 37: 293–329 (2009); Folke Tersman, "The reliability of moral intuitions: A challenge from neuroscience," *Australasian Journal of Philosophy* 86: 389–405 (2008).

not shot him since you would still have shot him, and the other way around. Similar remarks are true as regards Case Two.

However, even if it were true that other normative theories worth considering are partners in crime or have other just as serious difficulties, I still think that this defence would be somewhat problematic. The reason for that is that there is something to be said for the view that a consequentialist should be particularly disturbed by cases like Case One and Case Two. A consequentialist may of course be very disturbed by, for example, the fact that her view implies that is not wrong to kill one person in order to harvest his vital organs and transplant them to five other persons that would otherwise die. But she can at least say that the outcome is as good as it could be under those circumstances. However, a consequentialist that is disturbed by for example Case One cannot say that outcome is as good as it could be under the circumstances. It clearly is not: It would have been better if neither you nor I had shot Victim. Consequently, there seems to be a fairly strong reason why a consequentialist should be disturbed by cases like Case One and Case Two.¹⁴

4. Denying the Intuition

I now turn to the second strategy that the proponent of the Standard Version might employ to deal with Case One and Case Two. She might deny that we have the intuition that you and I respectively act wrongly in Case One and in Case Two.

Denying that we have the intuition that you and I respectively act wrongly in Case One and in Case Two might at a first glance seem desperate. Surely people are aware of what spontaneous judgments they make. Yes, but intuitions are not best thought of spontaneous judgements. It is better to think of them as judgments that are accepted not only in virtue of being implied by the person's moral principles.¹⁵ And it is not unlikely that people lose confidence in the belief that you and I respectively act wrongly in Case One and in Case Two when it becomes clear to them that neither your action nor my action individually makes any difference.

However, the proponent arguably still has to admit that something is not as it should be in these cases: It seems obvious that something is not *kosher* here. After all, Victim dies as a result of what you and I do, and it would have been better, if Victim had not been killed. If this is correct, the proponent of the Standard Version should accommodate this intuition, i.e., the more unspecific intuition that something is not as it should be in Case One and in Case Two.

¹⁴ Shelly Kagan essentially makes the same point. Shelly Kagan, "Do I Make a Difference?", *Philosophy & Public Affairs*, 39 (2011): 105-141, 107-108.

¹⁵ See Folke Tersman, "The reliability of moral intuitions: A challenge from neuroscience," 391.

As I said in the introduction to this chapter, there are primarily three suggestions in the literature of what it is that is not as it should be in Case One and Case Two: that you and I together act wrongly, that you and I act subjectively wrongly, and that you and I are uncooperative. As I also said in the introduction to this chapter, I shall argue that these three views are problematic.

5. Acting Wrongly Together

So, a proponent of the Standard Version may claim that although neither you nor I act wrongly in Case One and Case Two, you and I together act wrongly in these cases. More precisely, she may claim that you and I perform a group action in each case, and that these group actions are wrong. Both Frank Jackson and Torbjörn Tännsjö have (independently of one another) taken this view.¹⁶ However, this view faces some difficulties. The most important one is that it is not clear that you and I actually perform a group action in either of these cases. Both Jackson and Tännsjö have tried to handle this and other difficulties. In this section, I shall argue that although their attempts have something going for them, they are not entirely convincing.

In section 5.1, I briefly discuss the view that you and I together act wrongly in Case One and Case Two and note that it only works if you and I actually perform a group action in these cases. In section 5.2, I note that Tännsjö's attempt to handle this difficulty is problematic. In section 5.3, I note that Jackson's attempt also is problematic. Finally, in section 5.3, I also consider the view that you and I perform a plurality of actions in Case One and Case Two, and that these pluralities are wrong. (In chapter One, section 2.3, I distinguished between combinations of actions that are actions in their own right – for instance, group actions – and combinations of actions that are not actions in their own right. I referred to the latter category as pluralities of actions.) I shall argue that this suggestion is problematic too.

5.1 The Group Action Is Wrong

In Chapter Two at the end of section 2.3, I said that a combination of actions performed by a group of people is a group action just in case the members of the group jointly intend some aim. (I also said that many think that this is the case only if (a) each group member intends that they together pursue the aim, (b) each group member intends to do her part, (c) each group member forms

¹⁶ See Frank Jackson, 'Group Morality', 91–110 *Metaphysics and Morality: Essays in Honour of J.J. C. Smart*, ed. Philip Pettit, Richard Sylvan, and Jean Norman, Oxford: Basil Blackwell (1987). See also Torbjörn Tännsjö, "The Morality of Collective Actions," *The Philosophical Quarterly* 39: 221–228 (1989).

these intentions – at least to some extent – because she believes that the other members also do so, and (d) each member believes that (a), (b), and (c) are met and that the other members believe this too.) However, I have not said anything about whether you and I jointly intend to do anything in Case One and Case Two. Do we jointly intend to kill Victim in these cases?

Suppose we jointly intend to kill Victim in Case One. In order to increase our chances of killing Victim, we together decided that I shoot from one angle and you shoot from another, and so on. (It may of course still be true that each of us would have shot Victim regardless of the other.) In that case, it is (I shall assume) fair to say that you and I perform the group action “my shooting Victim, and your shooting Victim.” Furthermore, the Standard Version would then imply that this action is wrong. There is an alternative to this action whose outcome is intrinsically better. The outcome of the action, “my abstaining from shooting Victim, and your abstaining from shooting Victim”, is better since Victim would have survived if neither of us had shot him. I should perhaps stress that this does not imply that our individual actions are wrong. It does not imply, that is, that the action “my shooting Victim” is wrong and that the action “your shooting Victim” is wrong. It just implies that the action that consists of these two actions – i.e., the group action “my shooting Victim, and your shooting Victim” – is wrong.

This reply has some nice features. It addresses the challenge before us by pointing out that there is in fact an action that is wrong under the circumstances, something that may very well have been overlooked. Furthermore, it sharply distinguishes the moral status of individual actions from the moral status of compound actions and therefore avoids an important difficulty. As we shall see in Chapter Six (section 3.4), there are cases where a group action is wrong, but where the simple actions composing the group action are obligatory.

There is however a pretty obvious objection to this view. Two paragraphs back, I assumed that you and I jointly intended to kill Victim in Case One. But suppose now that we did not. The combination of actions, “my shooting Victim, and your shooting Victim,” is not under these circumstances an action. As I said in the very beginning of this section, in order for a combination of actions to count as a group action, the agents involved have to jointly intend some aim. In that case, the Standard Version does not apply and cannot therefore declare the combination of actions, “my shooting Victim, and your shooting Victim” wrong.

5.2 Tännsjö’s Reply

Tännsjö denies that you and I have to jointly intend to kill Victim in order for the combination of actions, “my shooting Victim, and your shooting Victim,” to count as a group action. Inspired by Donald Davidson – who says that “a man is the agent of an act if what he does can be described under an aspect

that makes it intentional” – Tännsjö holds the view that a group is the agent of a group action, if what the group does may somehow plausibly be described as being intentional.¹⁷

Consider again the combination “my shooting Victim, and your shooting Victim”. There is, according to Tännsjö, an intention corresponding to that combination: namely, the conjunction of my intention and your intention. That intention may, for example, be “*me wanting to kill Victim, and you wanting to kill Victim*”. (This is a slightly odd way to characterise the conjunctive intention. For example, my intention is rather *to kill Victim*, but I – and Tännsjö – put it like this to make the relation to the group action conspicuous.) This conjunctive intention is of course not a joint intention, but Tännsjö thinks that it is such that the combination, “my shooting Victim, and your shooting Victim” may plausibly be regarded as being intentionally performed by you and me. (That is to say, the group agent that is composed by you and me. According to Tännsjö, it is this entity that is the bearer of the conjunctive intention.)

As far as I can see, Tännsjö thinks that there is an intention corresponding to any combination of actions (where the actions are not performed by the same person) that is such that the combination in question may plausibly be regarded as being intentionally performed by the agents involved. For example, he notes that his “view of collective action may seem overly permissive. After all, on my account, a collective action such as my writing this essay and Brutus’ killing Caesar is allowed.”¹⁸

I think that Tännsjö’s view is problematic for several reasons. For one thing, (as just mentioned) it implies that countless combinations of actions that clearly do not seem to be actions come out as actions. It might not seem odd to claim that the combination of actions, “my shooting Victim, and your shooting Victim” in Case One is an action. The outcome is after all Victim’s death here. However, it is odd to claim that the combination “Tännsjö’s writing “The Morality of Collective Actions”, and Brutus’ killing Caesar” is an action in its own right.

Moreover, Tännsjö’s view implies that many combinations of actions you would not think of as having a moral status come out as having one. It might not seem odd to claim that the combination, “my shooting Victim, and your shooting Victim” in Case One is wrong. But it does seem odd to claim that the combination, “Tännsjö’s writing “The Morality of Collective Actions”, and Brutus’ killing Caesar” is wrong (or right for that matter). One reason why this might appear odd is that it in many cases it seems appropriate to punish those agents that intentionally act wrongly. However, you could punish the agent composed by Tännsjö and Brutus by punishing just Tännsjö. This is not an attractive implication.

¹⁷ Torbjörn Tännsjö, “The Morality of Collective Actions,” 226–228.

¹⁸ *Ibid.* 227.

Now, the fact that combination of actions that do not seem to have a moral status on Tännsjö's view come out as having one, also supports the claim that they are not actions. If they were actions, it would be difficult to explain why they would lack moral status. (Tännsjö does say – as regards the combination “Tännsjö's writing “The Morality of Collective Actions”, and Brutus' killing Caesar” – that although we “*may* assess the normative status of this action, if we please ... there is little *point* in doing so”.¹⁹ However, this claim does not seem to do anything to deny that it is *odd* that the combination under consideration comes out as wrong.)

Finally, many entities that are group agents on Tännsjö's view – for example, the group agent composed by Tännsjö and Brutus (associated with the combination, “Tännsjö's writing “The Morality of Collective Actions”, and Brutus' killing Caesar” – simply seem so different from us paradigmatic agents (i.e., us people) that it is difficult to believe that they are agents at all. Intuitively, we do not have proper parts that themselves are agents, for instance.

5.3 Jackson's Reply

Jackson also denies that you and I have to jointly intend to kill Victim in order for the combination of actions, “my shooting Victim, and your shooting Victim,” to be regarded as a group action. He thinks that it is sufficient that there in general sense is a common element – which does not have to be an intentional item – in what we do for the combination “my shooting Victim, and your shooting Victim,” to count as a group action. According to Jackson, the common element here is Victim's death. (Jackson's view is thus not a competitor to Tännsjö's view: They give different sufficient conditions.)

Jackson gives two examples to lend support to his view. He first invites us to consider Learner and Lowe, the writers of the musical *My Fair Lady*. They wrote the piece in close collaboration with one another: Learner wrote the lyrics and Lowe wrote the music. Now, Jackson wants us to suppose (counterfactually) that Learner and Lowe did not collaborate but produced their respective parts independently. (We also have to imagine that the parts somehow come together without their collaboration.) In that case, it would, according to Jackson, “still be true that they together wrote *My Fair Lady* ... without it being the case that they had a common goal, or indeed that one's action in any way affected the other's.”²⁰ The second example Jackson invites us to consider concerns two persons who – unaware of one another – build a tunnel through a mountain and, to their surprise, meet in the middle. Jackson rhetorically asks

¹⁹ Ibid. 227.

²⁰ Frank Jackson, “Group Morality,” 93.

us: “Isn’t opening a tunnel through the mountain a group action of theirs despite their failure to co-operate...?”²¹

But what shall we say in a case where there does not seem to be a readily distinguishable outcome corresponding to the combination in question? For example, there does not seem to be a readily distinguishable outcome corresponding to the combination, “my shooting Victim, and Fred Astaire’s dancing in the distance” in the case I discussed in the previous section. Jackson suggests that one may always come up with an outcome that has the right feel. He rhetorically asks: “But surely a common element can always be found – or cooked-up (grued-up)?”²²

I do not know whether this is true, but it might not be that important. Jackson could point out that there is a readily distinguishable outcome in Case One, i.e., Victim’s death. This also seems to be true of other cases like Case One. For example, each event – in a simple case of redundant causation – is sufficient for the same outcome. Jackson might therefore not need to claim that any combination of actions (involving more than one agent) is a group action to handle Case One and similar cases. It might be enough to claim that those combinations of actions where there is a clear common element, as for example those that have a readily distinguishable outcome, count as group actions.

I nevertheless think that it is better not to rely on Jackson’s view. Intuitively (and as I said in Chapter Two, section 2.3) actions – including group actions – are intentional. Of course, Jackson’s examples are intended to undermine this view. However, note that if we accept the view that there are actions that are not intentional it is difficult to deny that things – such as, for example, reflexes – are actions. (Reflexes also seem to have readily distinguishable outcomes.) But I do not think we would want to say that reflexes are wrong.

5.4 An Appeal to Pluralities of Actions

A proponent of the appeal to what we together do might concede that your action and my action in Case One and Case Two do not compose a group action but insist that this does not matter very much. She might claim instead that our actions belong to a plurality of actions that is wrong. More specifically she may adopt the following principle:

The Plurality Principle

A plurality of actions “ a_1, a_2, \dots, a_n ” is wrong, if and only if, there is an alternative to “ a_1, a_2, \dots, a_n ” whose outcome would be intrinsically better than those of “ a_1, a_2, \dots, a_n ”.

²¹ Ibid. 93.

²² Ibid. 93.

The Plurality Principle implies that the plurality “my shooting Victim, and your shooting Victim” is wrong. There is an alternative plurality whose outcome is better. The plurality “my abstaining from shooting Victim, and your abstaining from shooting Victim” is better since it would have been better if neither of us had shot Victim. (As I said in Chapter Two, section 2.3, some combinations of actions are pluralities of actions. Moreover, as I said in Chapter Two, section 2.4, the alternatives of a combination of actions are a function of its simple actions. For example, since you and I, respectively, could only either shoot or abstain from shooting, the alternatives to the combination “my shooting Victim, and your shooting Victim” are the following three: “my shooting Victim, and your abstaining from shooting Victim”, “my abstaining from shooting Victim, and your shooting Victim”, and “my abstaining from shooting Victim, and your abstaining from shooting Victim.”)

The Plurality Principle is also somewhat problematic, however. The trouble with this principle is that it does not seem entirely correct to attribute wrongness to things other than actions, and pluralities of actions are of course not actions in their own right. We do not attribute wrongness to natural events however devastating they may be. For example, we would not say that that hurricane Katrina and associated events were wrong. We may certainly say that they were bad, but hardly that they were wrong. Of course, it might not seem as odd to attribute wrongness to pluralities of actions, as it is to attribute wrongness to pluralities of natural events, but it is still (I think) somewhat odd. Moreover, just as Tännsjö’s view implies that countless combinations of actions that do not seem wrong come out as wrong, the Plurality Principle implies that countless pluralities of actions that do not seem wrong come out as wrong. For example, the plurality, “Tännsjö’s writing “The Morality of Collective Actions”, and Brutus’ killing Caesar” that I discussed above.

6. An Appeal to Subjective Moral Principles

A proponent of the Standard Version might also say that although neither you nor I act wrongly in Case One and Case Two, you and I, respectively, act *subjectively* wrongly in these cases. Subjective wrongness concerns the agent’s epistemic situation. And is it not reasonable to assume that it is such that you and I respectively have reason to believe – in both cases – that shooting Victim makes it likelier that Victim gets killed? In this section, I shall suggest that it might be better not to rely on this view.

In section 6.1, I explain how subjective wrongness might be understood, and in section 6.2, I suggest that it might not help us solve our problem.

6.1 Subjective Moral Principles

Philosophers often distinguish between *subjective* moral principles and *objective* moral principles. As I indicated in the previous section, according to subjective moral principles, the moral status of an action is partly determined by the agent's epistemic situation (i.e., the evidence available to the agent); whereas this is not so according to objective moral principles. The Standard Versions' three moral principles are objective moral principles. For example, the Standard Version's principle of wrongness says that an action is wrong if and only if it has an alternative whose outcome would be intrinsically better. As is easy to see, this principle does not consider the evidence available to the agent.

There are various ways to formulate a subjective counterpart to the Standard Version's principle of wrongness. For instance, it might be formulated as follows:

The Principle of Subjective Wrongness

An action *a* is *subjectively* wrong, if and only if, there is an alternative to *a* that has a higher expected value than *a*.

Why is this a subjective principle? Well, the expected value is calculated by multiplying the subjective probability (i.e., the probability reflecting the evidence available to the agent) of an action's possible outcomes with the intrinsic value of those outcomes, and summing up the results.²³ For example, suppose that I can join you on a rescue mission, and that the intrinsic value of each person's continued life is +1. Furthermore, suppose that the evidence available suggests that there are two possible outcomes associated with my action: Ten persons are saved, and no person is saved. Finally, suppose that the evidence available also suggest that I have one chance in two in saving ten persons and one chance in two in saving no person. In that case, the expected value of this option is +5 ($0.5 \times 10 + 0.5 \times 0$).

Just as the Standard Version's principle of wrongness has a subjective counterpart, the Standard Version's principle of obligatoriness and the Standard Version's principle of rightness have subjective counterparts too:

The Principle of Subjective Obligatoriness

An action *a* is *subjectively* obligatory, if and only if, *a* has a higher expected value than any of its alternatives.

²³ See for example Timmons on this. Mark Timmons, *Moral Theory: An Introduction*, Second Edition, Rowman & Littlefield Publishers, Inc. 2013, 132.

The Principle of Subjective Rightness

An action a is *subjectively* right, if and only if, there is no alternative to a that has a higher expected value.

Why have philosophers distinguished between objective principles and subjective principles? To see this, consider the following story: I could either join you on a rescue mission or I could go on my own rescue mission. I have no other option. The evidence available says that there are two possible outcomes associated with each action: Ten persons are saved, and no person is saved. Furthermore, the evidence available says that if I join you, I have nine chances in ten in saving ten persons and one chance in ten in saving no person. The evidence available also says that if I go off by myself, I have one chance in ten in saving ten persons and nine chances in ten in saving no person. I decide to join you on your rescue mission. Unfortunately, we fail spectacularly. No one is saved, but I would have saved ten if I had gone off by myself.

Now, according to the Standard Version I act wrongly since I have an alternative whose outcome would have been intrinsically better: ten more persons would have been saved, if I had gone off by myself. (As above, I assume that the intrinsic value of each person's continued life is +1.) However, would it not have been extremely irresponsible of me to do so? The evidence available said that it was a lot less likely that I would save ten persons if I go off by myself. To address this worry, the proponent of the Standard Version might appeal to subjective principles. She might claim that although I act wrongly, I act subjectively rightly. The expected value of joining you is +9 ($0.9 \times 10 + 0.1 \times 0$) and the expected value of going off by myself is +1 ($0.1 \times 10 + 0.9 \times 0$). (If I instead had gone off by myself she could have said that although I acted rightly, I acted subjectively wrongly.)

But how shall we understand the relation between objective moral principles and subjective moral principles? Well, the proponent of the Standard Version thinks that the objective moral principles provide the criteria that determine the wrongness, rightness, and obligatoriness of actions. For example, the Standard Version's principle of wrongness tells us what it is that makes an action wrong. The subjective principles, on the other hand, do not do that. Some proponents of the Standard Version think that subjective principles are guides to practical deliberation under conditions of uncertainty. They tell us what to do when – as often is the case – we do not know what the outcomes of our alternatives are. They tell me – in the case that I described in the previous paragraph – that I should join you. Other proponents of the Standard Version think that subjective principles guide the distribution of blame. For example, they think that a person is blameworthy if and only if she acts subjectively wrongly. They say that I – in the case that I described in the previous paragraph – am not blameworthy although I acted wrongly.

6.2 Subjective Moral Principles and Case One and Case Two

How can subjective moral principles be used to deal with Case One and Case Two?

Well, consider Case One, for example. Couldn't we say that although neither of us acts wrongly, each of us acts subjectively wrongly? It does not seem unreasonable to assume that the evidence available suggests that each of us only had some chance – say, one chance in two – in shooting a deadly shot, if he or she pulls the trigger. After all, few people are expert marksmen. Moreover, neither does it seem unreasonable to assume that the evidence available suggests that it is very unlikely that any of us causes any negative intrinsic value at all, if he abstains from pulling the trigger. After all, few people are exceptionally clumsy. In that case, the Principle of Subjective Wrongness implies that you and I, respectively, act subjectively wrong. The expected value of the outcome of what I did is clearly lower than the alternative. The same is of course true of you.

Well, the problem with saying that you and I respectively act subjectively wrongly in Case One is that I have not included in the description of Case One any information as regards the evidence available to the parties. It may therefore be the case that the evidence available suggests that each of us is almost certain to shoot a deadly shot in Case One. (If you find it difficult to believe this, imagine instead a thousand shooters.) Moreover, it may be the case that the evidence available suggests that there is some probability that each of us causes some negative intrinsic value, if he or she abstains from pulling the trigger. And if this is the case, the expected value of me shooting Victim might well be as high as me abstaining from shooting him, and the same is true of you. If so, my action is neither objectively nor subjectively wrong; and the same is true of your action.

Shelly Kagan has suggested that the agents in almost all realistic cases of over-determination act subjectively wrong, and that in those cases they do not, we just have to bite the bullet.²⁴ I do not think that this is a very good response for two reasons. First, we still have the intuition in the remaining cases that the agents involved act wrongly. Second, Kagans appeal to subjective wrongness would in any event be superfluous. It is also the case that the agents in almost all realistic cases of over-determination act objectively wrongly. It is very seldom true that people do not have an alternative whose outcome is intrinsically better.

²⁴ Shelly Kagan, "Do I Make a Difference?" *Philosophy & Public Affairs* 39 (2011): 105-141, pp. 107-108. See also Alastair Norcross, "Harming in Context," *Philosophical Studies* 123 (2005): 149-173.

7. An Appeal to Uncooperativeness

Finally, a proponent of the Standard Version might say that although neither you nor I act wrongly in Case One and Case Two, you and I, respectively, are uncooperative in the sense that each of us would still have killed Victim even if the other had acted differently or offered to act differently.²⁵ Neither of us would have agreed to bring about the best outcome we could bring about. Some proponents of the Standard Version think that you and I – by being uncooperative in the stated sense – display a bad character.

There is one problem with this response, though. Consider Case Five again:

Case Five

Case Five is exactly like Case Two except that my shot counterfactually depends on your shot. (I am simply the sort of person that wouldn't shoot someone unless someone else goes first.) We may imagine that you did shoot first, but my bullet nevertheless reached Victim before yours.

As I said when I discussed Case Five above, it too challenges the Standard Version. My action is not wrong here, according to the Standard Version, since I do not have an alternative whose outcome is intrinsically better: Victim would still have been killed by you, if I had not shot him. Intuitively, I do act wrongly, however. If I act wrongly in Case Two where my action does not depend on yours, it is difficult to see why I would not be action wrongly in Case Five where my action does depend on yours (perhaps particularly so when my character seems shady). But the appeal to uncooperativeness does not work here since I am not uncooperative in the stated sense.

²⁵ Frank Jackson, Felix Pinkert, and Michael Zimmerman have made suggestions along these lines. See Frank Jackson, "Which Effects?", 50; Felix Pinkert, "What If I Cannot Make a Difference (and Know it)?", *Ethics* 125: 971–998 (2015); Michael J. Zimmerman, *The Concept of Moral Obligation*, Cambridge: Cambridge University Press (1996), 266–268.

Chapter Four, Causal Consequences

1. Introduction

In the previous chapters I have at several occasions noted that the Standard Version of Act-Consequentialism implies that neither you nor I act wrongly in Case One and Case Two (i.e., the cases where you and I shoot and kill a third person called “Victim”). Furthermore, in the previous chapter I argued that both of the two strategies that the proponent of the Standard Version might employ to handle Case One and Case Two (i.e., accepting that the implications of Standard Version in Case One and Case Two are counterintuitive but explain away the intuitions, and denying that these implications are counterintuitive) are problematic.

In this chapter I shall discuss a version of Act-Consequentialism that might seem to fare better. The Standard Version interprets the term “outcome of an action” as referring to the possible world that would be the case, if the action were performed. The version I have in mind, the “Causal Consequences Version of Act-Consequentialism”, understands “outcome of an action” as instead referring to the complete causal consequences of the action that would be the case, if the action were performed. And is it not clear, for example, that you and I, respectively, cause the state of affairs that Victim dies at t_1 in Case One (i.e., that this state is a causal consequence of each of our actions in Case One)?²⁶ I shall reject the Causal Consequences Version, however.

I shall proceed as follows: In section 2, I introduce the Causal Consequence Version. I explain why it might be thought to handle Case One and Case Two. In section 3, I discuss a difficulty for the Causal Consequence Version. According to a very natural view on causation – i.e., that causal consequences counterfactually depend on their causes – the state of affairs that Victim dies at t_1 is not a causal consequence of my action and your action respectively in Case One. In section 4, I consider an account of causation that does have the desired implications in Case One. In section 5, I argue that we should reject the Causal Consequence Version regardless of whether the account of causation I discuss in section 4 is plausible. The Causal Consequence Version has some very implausible implications.

²⁶ Note that I relate states of affairs (i.e., that Victim dies at t_1) and events (i.e., my action and your action). The standard view is that causal relata are events. I explain why in the beginning of Section Two.

2. The Causal Consequence Version

The Causal Consequence Version of Act-Consequentialism is of course a version of Act-Consequentialism. I presented Act-Consequentialism in Chapter Two (section 2), but here it is again:

Act-Consequentialism

An action *a* is *wrong*, if and only if, there is an alternative to *a* whose outcome would be intrinsically better than that of *a*.

An action *a* is *obligatory*, if and only if, the outcome of *a* is intrinsically better than the outcome of any of its alternatives would be.

An action *a* is *right*, if and only if, there is no alternative to *a* whose outcome would be intrinsically better than that of *a*.

As I said in Chapter Two, the Standard Version makes two claims that set it apart from other versions of Act-Consequentialism. First, it claims that the intrinsic value of an outcome is determined by well-being. Second, it claims that outcome of an action is the entire possible world that would be the case, if the action were performed.

The Causal Consequence Version agrees with the Standard Version that the intrinsic value of an outcome is determined by welfare or well-being. However, it says that the outcome of an action is the complete causal consequence of the action that would be the case, if the action were performed rather than (as the Standard Version claims) the entire possible world that would be the case, if the action were performed.

The Causal Consequence Version is concerned with the *complete* causal consequences of an action. It is easy to see why. Suppose that Gertrude kills Victim with one shot. Moreover, suppose that Gertrude also saves Gertrude's sister and Gertrude's mother with that same shot. In that case, Gertrude plausibly causes the state of affairs that Victim dies at t_1 . However, to properly assess Gertrude's action, we obviously have to take into account the fact that she also causes the state of affairs that Gertrude's sister is saved and the state of affairs that Gertrude's mother is saved.

But how do we specify the complete causal consequences of an action? I shall say that a state of affairs *S* is a *simple* causal consequence of an event *c* if and only if (1) *S* is a simple state of affairs and (2) *S* is a causal consequence of *c*. Furthermore, I shall say that a that a state of affairs *S* is a *simple* state of affairs if and only if it does not have the logical form of a conjunction, disjunction, implication or equivalence. (A simple state of affairs may thus have the form of negation.) Now, we are in the position to say what the complete

causal consequences of an action are: A state of affairs S is the *complete* causal consequence of an event c , if and only if, S includes all, and only, simple causal consequences of c .²⁷

2.1 Case One and Case Two

But why do the proponents of the Causal Consequence Version think that it fares better in Case One and Case Two than the Standard Version? (Of course, there may be proponents of the Causal Consequence Version that do not think that it does better in Case One and Case Two than the Standard Version. These proponents might think that a theory may imply either that neither you nor I act wrongly in Case One, or that both you and I do so, and embrace the Causal Consequence Version for some other reasons. But let us – for the sake of expediency – speak as if all proponents of the Causal Consequence Version think that it does.) Well, the proponents of the Causal Consequence Version think that it has the intuitively correct implications in Case One and in Case Two.

Consider first Case One:

Case One

You and I shoot and kill Victim. Our bullets arrive at the same time and each shot would have killed Victim by itself. Victim would not have been killed, if neither of us had pulled the trigger. Apart from shooting Victim, you and I have only one alternative each: each could also have abstained from shooting Victim. Furthermore, each of us would have performed his or her alternative – i.e., each of us would have abstained from shooting Victim – if he or she had not shot him. We act independently of one another, Victim's level of well-being is not affected by either of the shots taken by itself (e.g., Victim does not experience any pain), and no one else is affected. Finally, Victim would have gone on to live a life well worth living, if neither you nor I had shot him.

The proponents of the Causal Consequence Version think it is pretty clear that you and I respectively are *causing* bad things to be the case in Case One. They think that it is clear that you and I respectively cause the state of affairs that Victim dies at t_1 in Case One, or what amounts to the same, that the state of affairs that Victim dies at t_1 is a causal consequence of your action and my action respectively in Case One. (Clearly, c is a cause of S (or c causes S) if and only if, S is a causal consequence of c .)

This is not an unreasonable view. Consider a case that is precisely as Case One except that your gun does not fire. In this case, my action is clearly a

²⁷ See Erik Carlson, *Consequentialism Reconsidered*, Dordrecht, Kluwer Academic Publishers (1995), 48–58.

cause of the state of affairs that Victim dies at t_1 . If so, it seems plausible to say that my action in Case One should also qualify as a cause. The difference between the case where I shoot Victim alone and Case One is that you also shoot Victim in Case One. However, it would be slightly odd to say that this “addition” *deprives* my action of its causal powers. How could it do that? It does not interfere with my action or its causal process.

Consider now Case Two:

Case Two

You and I shoot one shot each at Victim. My bullet arrives first and kills Victim. Your bullet arrives very shortly after mine but does not kill Victim since my bullet has already done that. Your bullet would however have killed Victim, if my bullet had not done so first. Victim would not have been killed, if neither of us had pulled the trigger. Apart from shooting Victim, you and I have only one alternative each: each could also have abstained from shooting Victim. Furthermore, each of us would have performed his or her alternative – i.e., each of us would have abstained from shooting Victim – if he or she had not shot him. We act independently of one another, Victim’s level of well-being is not affected by my shot (e.g., Victim does not experience any pain), and no one else is affected. Finally, Victim would have gone on to live a life well worth living, if neither you nor I had shot him, and the extra time Victim would have lived, if I had not shot him would neither have been intrinsically good nor intrinsically bad for him.

The proponents of the Causal Consequence Version think it is clear that I but not you *cause* the state of affairs that Victim dies at t_1 in Case Two. I do not think anyone would dispute that.

3. A Difficulty for The Causal Consequence Version

As I said, according to a very natural view on causation (i.e., that causal consequences counterfactually depend on their causes), the state of affairs that Victim dies at t_1 is *not* a causal consequence of my action and your action, respectively, in Case One. Therefore, the proponent of the Causal Consequence Version cannot just assume that the state of affairs that Victim dies at t_1 is a causal consequence of my action and your action respectively in Case One (although it might seem to be a reasonable view). However, unless the state of affairs that Victim dies at t_1 is a causal consequence of my action and your action respectively in Case One, it is difficult to see how the Causal Consequence Version can imply that you and I respectively act wrongly here.

This section is organised as follows: In section 3.1, I introduce the very natural account – the “Simple Counterfactual Account of Causation” – I have mentioned and explain why one might think that it is plausible. In section 3.2, I show that the Simple Account implies that neither you nor I cause the state of affairs that Victim dies at t_1 in Case One. Here I also consider some replies on behalf of the Causal Consequence Version. I suggest that they are not very persuasive.

3.1 The Simple Counterfactual Account of Causation

As I indicated, a very natural view (popularised by David Lewis) is that that causation should be analysed in terms of counterfactual dependence.²⁸ (An event e counterfactually depends on an event c , if and only if, e would not have occurred, if c had not occurred.) The Simple Counterfactual Account of Causation expresses this idea:

The Simple Counterfactual Account of Causation

A state of affairs S is a *causal consequence* of an event e , if and only if,

- (1) S obtains
- (2) S would not have obtained, if e had not occurred.

The Simple Account is intuitively plausible. It handles both “ordinary” cases of causation and problematic cases like Case Two. Consider first an “ordinary” case: Suppose I throw a brick through a window and that the window shatters. That the window shatters certainly seem to be a causal consequence of my throw, and this is exactly what the Simple Account implies: Clause (1) is satisfied since the window shatters, and clause (2) is satisfied because the window would not have shattered, if I had not thrown the brick.

Consider now Case Two (i.e., the case where both you and I shoot, but where you don’t kill Victim). The state of affairs that Victim dies at t_1 in Case Two is a causal consequence of my action, according to the Simple Account: Clause (1) is satisfied since the state of affairs that Victim dies at t_1 obtains, and clause (2) is satisfied since that state of affairs would not have obtained, if I had not shot him. Of course, Victim would still have been killed, but he would have been killed at a later time.

The Simple Account relates an event and a state of affairs. (Most other accounts relate events, but the Simple Account suits our purposes here since the Standard Version’s view on outcomes – i.e., that the outcome of an action is the possible world that would be the case, if the action were performed – also relates events and states of affairs since actions are events and it conceives of possible worlds as a certain type of states of affairs. It might be thought that it

²⁸ David Lewis, “Causation”, *Journal of Philosophy* 70: 556–567, 1973.

is problematic to take an event and a state of affairs as relata of a causal relation. Events are concrete things located in space and time, whereas states of affairs are abstract beings located outside space and time; but it seems that only concrete things located in space and time can interact properly, that – in a manner of speaking – only concrete things can push and shove, and be pushed and be shoved.

I do not think that it is clear that the relata of a causal relation must be concrete things located in space and time. Plenty of other philosophers agree. For example, Jonathan Bennett (who advocates fact causation) claims that it is not the relata of a causal relation that do the pushing and the shoving, that is rather done by physical objects:

I grant that facts cannot behave like elbows in the ribs, but we know what items do play that role — namely, elbows. In our world the pushing and shoving and forcing are done by *things* — elementary particles and aggregates of them — and not by any relata of the causal relation.²⁹

3.2 Case One

According to the Simple Account, the state of affairs that Victim dies at t_1 is neither a causal consequence of my action nor of your action in Case One (i.e., the case where you and I shoot Victim and our bullets arrive at the same time).

Consider my action. Clause (1) is satisfied since the state of affairs that Victim dies at t_1 obtains. However, clause (2) is not satisfied since that state of affairs would still have obtained, if I had not shot him, since you would have killed him at that time. Similar things are obviously true of your action. Consequently, the Simple Account fails to lend support to the claim that the state of affairs that Victim dies at t_1 is a causal consequence of my action and your action, respectively, in Case One.

I cautioned above (in the beginning of section 3) that the proponents of the Causal Consequence Version cannot just assume that the state of affairs that Victim dies at t_1 is a causal consequence of my action and your action respectively in Case One, because the Simple Account implies that it is not. The proponents of the Causal Consequence Version might protest that it is, in fact, intuitively clear that the state of affairs that Victim dies at t_1 is a causal consequence of each of our actions in Case One. (They need no account of causation to corroborate that!) They might claim that the fact that the Simple Account implies that the state of affair that Victim dies at t_1 is not a causal consequence of my action and your action respectively in Case One simply shows that the Simple Account is not as plausible as one might initially think.

I have some sympathy for this reply. However, it does have a somewhat limited appeal. There are people that think that our intuitions in cases like Case

²⁹ Jonathan Bennett, *Events and their Names*, Indianapolis, Hackett Publishers (1988), p. 22.

One, i.e., in cases of over-determination are unclear. (A case of redundant causation is a case of over-determination if and only if either both redundant causes are causes or neither of them is a cause, and the causal process running from each redundant cause proceeds without being interrupted.)³⁰ For example, consider Lewis's view on cases of over-determination (or *symmetrical* cases as he also calls them):

Some cases of redundant causation are symmetrical: Both candidates have an equal claim to be called causes of the effect. Nothing, either obvious or hidden, breaks the tie between them. It may be unclear whether we ought to say that each is a cause or whether we ought to say that neither is a cause (in which case we can still say that the combination of the two is a cause). But anyway it is out of the question to say that one is a cause and the other isn't. Because it's unclear what we want to say, these symmetrical cases are not good test cases for proposed analyses of causation. Set them aside.³¹

So, according to Lewis, it is not clear that you and I, respectively, cause Victim's death in Case One (or – what amounts to the same – that Victim's death is a causal consequence of either of our actions).

I also claimed above (also in the beginning of section 3) that it is difficult to see how the Causal Consequence Version might imply that you and I, respectively, act wrongly in Case One, unless it is true that that the state of affairs that Victim dies at t_1 is a causal consequence of your action and my action, respectively, in Case One. The proponents of the Causal Consequence Version might want to question this claim. They might want to point out that there are other states of affairs that do the job in Case One. For example, Victim would in all likelihood have died a slightly different death, if, for instance, I had not shot him, since his body would then have been pierced by one bullet instead of two. So, the state of affairs that Victim dies at t_1 in some particular manner is probably a causal consequence of my action in Case One, according to the Simple Account. Consequently, the proponents of the Causal Consequence Version might point out that the fact that the Simple Account implies that the state of affairs that Victim dies at t_1 is not a causal consequence of each of our actions, does not seem to damage the Causal Consequence Version after all.

This is not a very good reply, however. We may imagine cases very similar to Case One where Victim would not have died a slightly different death, if one of us had refrained from shooting him. For example, suppose you and I shoot Victim, but our bullets do not hit Victim directly. Instead, they arrive simultaneously at a mechanism that is then activated and kills Victim. Furthermore, each shot is sufficient by itself to activate this mechanism and it

³⁰ See Chapter Two for a discussion on this.

³¹ David Lewis "Causation as Influence," 74–106 in John Collins, Ned Hall, L. A. Paul (ed.), *Causation and Counterfactuals*, Cambridge Massachusetts, MIT Press 2004, 80.

would not have been activated, if neither you nor I had hit it with our respective bullets. Here, Victim would not have died a slightly different death if I had not fired my gun. (I am assuming that the mechanism has the same effect on Victim regardless of whether it is activated by one bullet or two bullets.) But we would want to say in this case too that you and I respectively act wrongly.

4. A Helping Hand

In this section I shall try to give the proponents of the Causal Consequence Version a helping hand. I shall consider an account of causation that does imply that the state of affairs that Victim dies at t_1 is a causal consequence of your action and my action, respectively, in Case One. As I have said several times, it is difficult to see how the Causal Consequence Version might imply that you and I respectively act wrongly in Case One unless this is the case.

This is my plan: in section 4.1, I point out that the Simple Account is probably false. In section 4.2, I present the account of causation that does have the desired implications. I want to stress already at this point that I do not regard this account as an analysis of causation. As we shall see, such an analysis might well be circular. I am merely trying to help the proponents of the Causal Consequence Version. In sections 4.3–4.5, I discuss some cases that this account, which I simply call “Causation”, handles.

4.1 An Objection to the Simple Account

Before I discuss Causation, I should point out that although the Simple Account initially might seem plausible, it is probably false.

Consider again Case Two (i.e., the case where you and I shoot Victim but my bullet arrives first). I noted that the state of affairs that Victim dies at t_1 in Case Two is a causal consequence of my action, according to the Simple Account (and regarded that as a reason in favour of the Simple Account). However, there are other problematic states of affairs in the vicinity. For example, the state of affairs that Victim dies *prematurely* is not a causal consequence of my action, according to the Simple Account. Victim would have died prematurely even if I had not shot him since you would have shot him only a moment later. But intuitively, the state of affairs that Victim dies prematurely is a causal consequence of my action. Consequently, the Simple Account does not, after all, have the correct implications in Case Two.

4.2 Causation

As indicated, there is an account of causation – which I call “Causation” – that does imply that the state of affairs that Victim dies at t_1 is a causal consequence of your action and my action, respectively, in Case One.

Here it is:

Causation

A state of affairs S is a *causal consequence* of an event c , if and only if, c is a redundant cause of S , and there is an uninterrupted causal process that runs from c to S .

In Chapter Two, I specified under what circumstances an event is a redundant cause of another event. Now I need to say under what circumstances an event is a redundant cause of a state of affairs since Causation understands causal consequences in terms of states of affairs. (What I say here is very similar to what I say there, however.) Under what circumstances, then, is an event a redundant cause of a state of affairs?

I shall make the following claim:

Redundant Causation

An event c is a *redundant cause* of a state of affairs S , if and only if, S obtains and there are events $c_1 \dots c_n$ such that

- (1) S would have obtained, if c had not occurred and $c_1 \dots c_n$ had occurred
- (2) S would have obtained, if none of $c_1 \dots c_n$ had occurred and c had occurred
- (3) S would not have obtained, if neither c nor none of $c_1 \dots c_n$ had occurred.

Now, the state of affairs that Victim dies at t_1 is a causal consequence of each of our actions in Case One, according to Causation. Note first that our actions are redundant causes of the state of affairs that Victim dies at t_1 in Case One, according to Redundant Causation. Consider your action, for instance. Clause (1) is satisfied because the state of affairs that Victim dies at t_1 would have obtained, if you had not shot Victim and I had since one shot is enough to kill him. (The variables $c_1 \dots c_n$ may pick out only one event.) Likewise, clause (2) is satisfied because the state of affairs that Victim dies at t_1 would have obtained, if I had not shot Victim and you had done so since one shot is sufficient to kill him. Finally, clause (3) is satisfied since the state of affairs that Victim dies at t_1 would *not* have obtained, if neither you nor I had shot Victim since it takes at least one of our shots to kill him. Consequently, your action is a redundant cause of the state of affairs that Victim dies at t_1 , according to Redundant Causation. Similar claims are true of my action.

Note furthermore that the causal process that runs from your action to the state of affairs that Victim dies at t_1 is not interrupted. This is also true of the causal process that runs from my action to the state of affairs that Victim dies at t_1 .

But what is a causal process? Well, I am not going to provide an analysis of causal processes here. I think we have an intuitive understanding of these processes. I am confident that in many cases we will recognise a causal process when we see one. It is clear that we are able to do so in Case One and Case Two, for instance. Furthermore, as I have stressed, I do not intend Causation to provide an analysis of causation. If it were an analysis of causation, it might well be circular since it seems plausible that the notion of a causal process ultimately rests on the notion of a cause.

4.3 Causation and Case Two

Causation has the correct implications in Case Two (i.e., the case where you and I shoot Victim but my bullet arrives first). That is to say, it implies both that the state of affairs that Victim dies at t_1 and that the state of affairs that Victim dies prematurely is a causal consequence of my action, but not of yours.

Consider first the state of affairs that Victim dies at t_1 . Although it might not be immediately obvious, my action is a redundant cause of that state of affairs, according to Redundant Causation. But how can that be? Clause (1) might not seem to be satisfied. It is not the case that Victim would have died at t_1 , if I had not shot him and you had since your shot would have killed him a bit later.

However, let the variables $c_1 \dots c_n$ pick out my action again (the variable c picks out my action) instead of yours. The antecedents of the counterfactuals in clause (1) and clause (2) would in that case be contradictions – i.e., if my shooting Victim had not occurred and my shooting Victim had occurred; and one may claim – as some do, for example, David Lewis – that the counterfactuals in clause (1) and clause (2) are true on account of that.³² Finally, clause (3) would also be satisfied. It is true that the state of affairs that Victim dies at t_1 would not have obtained, if neither “my shooting Victim” nor “my shooting Victim” had occurred (i.e., if “my shooting Victim” had not occurred). Consequently, under the assumption that counterfactuals with contradictions as antecedents are true (which I endorse), my action is a redundant cause here.

Furthermore, the causal process between my action and the state of affairs that Victim dies at t_1 is not interrupted. Nothing gets in between my action and this state of affairs.

What about your action? It is enough to notice that it does not satisfy clause (3) of Redundant Causation: It is not true that the state of affairs that Victim

³² David Lewis, “Causation,” *Journal of Philosophy* 70: 556–567, 1973, p. 560.

dies at t_I would not have obtained, if you had not shot Victim, since I would have shot him. Consequently, this state of affairs is not a causal consequence of your action, which obviously is the correct implication.

Consider now the state of affairs that Victim dies prematurely. I noted that it is not a causal consequence of my action, according to the Simple Account. However, my action is a redundant cause of that state of affairs. Clause (1) is satisfied because the state of affairs that Victim dies prematurely would have obtained, if I had not shot Victim and you had: one shot is sufficient to kill Victim, and the fact that you would have killed Victim a little later obviously does not make it false that he would have died prematurely. Moreover, clause (2) is satisfied: the state of affairs that Victim dies prematurely would obviously have obtained, if I you had not shot Victim and I had done so. Finally, clause (3) is satisfied since the state of affairs that Victim dies prematurely would *not* have obtained, if neither you nor I had shot Victim: he would have gone on to live a long life well worth living, if you and I had not attacked him.

Furthermore, the causal process between my action and the state of affairs that Victim dies prematurely is not interrupted. Again, nothing gets in between my action and this state of affairs.

What about your action? Well, it is enough to note that the causal process between your action and the state of affairs that Victim dies prematurely is interrupted. It is interrupted by the causal process running from my action. Your bullet reaches Victim when he is dead.

4.4 Causation and the “Ordinary” Case

Causation also has the correct implication in the “ordinary” case I discussed in section 3.1 (i.e., the case where I alone throw a brick through a window that then shatters). That is to say, the state of affairs that the window shatters is a causal consequence of my throw, according to Causation. It is clear that there is an uninterrupted causal process that runs from my throw to the state of affairs that the window shatters. Consequently, it is enough that I show that my throw is a redundant cause of the state of affairs that the window shatters, according to Causation.

Let the variable c pick out my action, and let the variables $c_1 \dots c_n$ pick out my action again. Clause (1) and clause (2) are in that case satisfied since the antecedents would be contradictions (i.e., if my throw had not occurred and my throw had occurred). Clause (3) would also be satisfied since it is true that the state of affairs that the window shatters would not have obtained, if neither “my throw” nor “my throw” had occurred. That is to say, the state of affairs that the window shatters would not have obtained, if my throw had not occurred.

Before I move on there is one further issue I would like to discuss. In section 3.1, I showed that my throw (in the case we are considering) is a cause of the state of affairs that the window shatters, according to Simple Account.

Moreover, we just saw that my throw is redundant cause of the state of affairs that the window shatters, according to Redundant Causation. This suggests that whenever an event is a cause of a state of affairs, according to the Simple Account, this event is also a redundant cause of the same state of affairs, according to Redundant Causation. (Or whenever a state of affairs is a causal consequence of an event, according to the Simple Account, that event is a redundant cause of that state of affairs, according to Redundant Causation.)

But is this true? Yes, since, as we saw two paragraphs back, when the variables c and $c_1...c_n$ of Redundant Causation pick out the same event (there, my throw), clause (1) and clause (2) of Redundant Causation are automatically satisfied (their antecedents are then contradictions and they are therefore vacuously true), and since clause (3) of Redundant Causation then is equivalent to the *analysans* of the Simple Account.

This means that if we want to find out whether an action is a redundant cause, according to Redundant Causation, it is sometimes sufficient to establish that it is a cause, according to the Simple Account. Furthermore, this also means that if we want to find out whether an action is a cause, according to Causation, it is sometimes sufficient to establish that it is a cause, according to the Simple Account, provided that we also establish that there is an uninterrupted causal process that runs from the action in question to the relevant state of affairs.

4.5 Causation and Redundant Causes

One might wonder why Causation invokes redundant causes? To establish that an event is a cause of some state of affairs, is it not enough to establish that there is an uninterrupted causal process running from the event to the state of affairs? For example, to establish that my action and your action, respectively, are causes of the state of affairs that Victim dies at t_1 in Case One, is it not enough to establish that there is an uninterrupted causal process running from each of our actions to this state of affairs? Well, I don't think it is.

Consider this case: There is a runaway trolley heading towards Victim. It will kill him, if it hits him. Further down the track there is a lever that – if pulled – will divert the trolley on to a side track away from Victim. However, further down the side track there is a yet further lever that – if pulled – will direct the trolley back on to the main track towards Victim again. You divert the trolley on to the side track, by pulling the first lever. But I redirect the trolley on to the main track towards Victim again, by pulling the second lever.

Intuitively, you do not cause that Victim dies at t_1 . But the causal process between your action and the state of affairs that Victim dies at t_1 does not seem to be interrupted.

However, your action is not a redundant cause, according to Redundant Causation: Clause (2) is not satisfied since the state of affairs that Victim dies at t_1 would not have obtained if you had pulled the lever and I had not since

the trolley would not then have returned to the main track. Furthermore, clause (3) is not satisfied since the state of affairs that Victim dies at t_1 would still have obtained if neither you nor I had pulled a lever since the trolley would not then have been diverted to the side track but instead continued towards Victim. (If we let the variables $c_1 \dots c_n$ pick out your action again instead of my action, clause (3) would not be satisfied. The state of affairs that Victim dies at t_1 would have obtained, if neither “your pulling the lever” nor “your pulling the lever” had occurred; i.e., that Victim dies at t_1 would have obtained even if you had not pulled the lever.)

Furthermore, the proponents of the Causal Consequence Version would presumably not want to say that you cause the state of affairs that Victim dies at t_1 in this case since they would not want to say that you act wrongly here. After all, you try to save Victim by redirecting the trolley away from him, and you would succeed, if it were not for me.

5. Rejecting the Causal Consequence Version

In this section, I shall argue that the Causal Consequence Version should be rejected even if it were true that the state of affairs that Victim dies at t_1 is a causal consequence of your action and my action, respectively, in Case One. There are two reasons for that. First, the Causal Consequence Version implies that you and I respectively act wrongly in Case One and that you alone act wrongly in Case Two, only if we assume that the state of affairs that Victim dies at t_1 is intrinsically bad. But this is an implausible assumption. Second, the Causal Consequence Version has implausible implications.

I shall proceed as follows. In section 5.1, I show that the Causal Consequence Version handles ordinary cases. In section 5.2, I explain why we have to assume that the state of affairs that Victim dies at t_1 is intrinsically bad. I also explain why we should hesitate to make this assumption. I also show that the Causal Consequence Version has (if we make the assumption) some undesirable implications.

5.1 The Causal Consequence Version and Ordinary Cases

Let us see how the Causal Consequence Version works in an ordinary case.

Suppose that I have two alternatives at t_1 : I could shoot and kill Victim or I could stand still. Suppose I actually shoot and kill him. What are the complete causal consequences of me doing that? Well, it depends on what happens in the closest possible world where I do not shoot and kill him. Suppose that possible world is the possible world where I stand still. Furthermore, suppose that Victim does not die at t_1 in that possible world, but that he instead enjoys an ice cream at t_2 (among other things). In that case, the complete causal con-

sequence of me shooting Victim includes at least the state of affairs that Victim dies at t_1 , according to Causation. The state of affairs that Victim dies at t_1 is a causal consequence of my shot, according to the Simple Account, since Victim would not have died at t_1 , if I had not shot him. And as I noted in section 4.4, whenever a state of affairs is a causal consequence of an event, according to the Simple Account, this event is a redundant cause of that state of affairs, according to Redundant Causation. Furthermore, it is evident that there is an uninterrupted causal process that runs from my shot to the state of affairs that Victim dies at t_1 .

What would the complete causal consequences be of me standing still? Well, consider the world where I stand still (i.e., the closest possible world to the actual world where I do not shoot and kill Victim). The complete causal consequences of me standing still depends on what happens in the closest possible world where I do not stand still. Suppose that this is the possible world where I shoot and kill Victim. In that possible world Victim dies at t_1 and does not therefore enjoy an ice cream at t_2 . So, the complete causal consequences of me standing still would at least include the state of affairs that Victim enjoys an ice cream at t_2 , according to Causation. The state of affairs that Victim enjoys an ice cream at t_2 , is a causal consequence of me standing still, according to the Simple Account, since Victim would not have enjoyed an ice cream at t_2 , if I had not stood still. Again, as I noted in section 4.4, whenever a state of affairs is a causal consequence of an event, according to the Simple Account, that event is a redundant cause of that state of affairs, according to Redundant Causation. Moreover, there seems to be an uninterrupted causal process that runs from me standing still to the state of affairs that Victim enjoys an ice cream at t_2 .

Now, suppose that the net intrinsic value of the complete causal consequences of me shooting Victim is zero. This is a plausible claim (in the light of the assumptions I made above). The state of affairs that Victim dies at t_1 is arguably neither intrinsically good nor intrinsically bad for him. Other relevant states of affairs are arguably also neither intrinsically bad nor intrinsically good for Victim since I kill him. Furthermore, suppose that the net intrinsic value of the complete causal consequences me standing still would be positive. This is also a plausible assumption. The state of affairs that Victim enjoys an ice cream at t_2 is arguably intrinsically good for him, and we may assume that the net intrinsic value of all other states of affairs included in the complete causal consequence of me standing still would be positive, or at least zero. In that case, the Causal Consequence Version implies that I act wrongly by shooting Victim since I have an alternative to what I do whose outcome would be intrinsically better.

But let us complicate matters somewhat. Return to the case in the previous paragraph but suppose that I had yet a further alternative. Suppose that apart from the alternatives mentioned above I could also have performed a little dance. Would the complete causal consequences of me standing still under

these circumstances still include the state of affairs that Victim enjoys an ice cream at t_2 ? Well, that depends on what I do in the closest possible world where I do not stand still. Suppose this time that this possible world is not the world where I shoot and kill Victim (as it was in the case in the previous two paragraphs) but that it is the possible world where I perform a little dance. In that case, it might seem that the causal consequences of me standing still would *not* include this state of affairs. The state of affairs that Victim enjoys an ice cream at t_2 is *not* a causal consequence of me standing still here, according to the Simple Account because Victim would still be enjoying an ice cream at t_2 under these circumstances.

But would not the causal consequences of me performing a little dance include the state of affairs that Victim enjoys an ice cream at t_2 ? Well, consider the possible world where I perform the little dance. The causal consequences of me performing the little dance depend on what I do in the closest possible world where I do not perform the little dance. Suppose that world is the world where I stand still. In that case, it might seem that the causal consequences of me performing a dance would not include this state of affairs. The state of affairs that Victim enjoys an ice cream at t_2 is *not* a causal consequence of me performing a little dance here, according to the Simple Account because Victim would still be enjoying an ice cream at t_2 under these circumstances. However, if this is the case, the Causal Consequence Version implies that it is right to shoot Victim. The net intrinsic value of the causal consequences of me standing still would be zero. The same would be true of the causal consequences of me performing a little dance. (I am still assuming that the closest world to the world where I shoot and kill Victim is the possible world where I stand still.)

However, the Causal Consequence Version does not imply that it is right to shoot Victim: Me standing still and me performing a little dance would (perhaps contrary to first impression) both be redundant causes of the state of affairs that Victim enjoys an ice cream at t_2 . Consider the possible world where I stand still. The variable c naturally picks out me standing still. But which action (or actions) should the variables $c_1 \dots c_n$ pick out? We cannot let them pick out me performing a little dance. In this possible world (i.e., where I stand still) I do not also perform a little dance.

But we might let these variables pick out an event further back – for example, my decision that if I do not stand still I will instead perform a little dance. Clause (1) would, then, be satisfied because the state of affairs that Victim enjoys an ice cream at t_2 would have obtained, if I had not stood still and taken the decision (i.e., if I do not stand still, I shall perform a little dance) because I would then surely have performed a little dance. Clause (2) would also be satisfied, under these circumstances, since the state of affairs that Victim enjoys an ice cream at t_2 would have obtained, if I had not taken the decision and had stood still. Finally, the state of affairs that Victim enjoys an ice cream at t_2 would *not* have obtained, if I had performed neither of these actions because

I would then have shot Victim. So, the Causal Consequence Version does after all handle this more complicated case.

5.2 The Intrinsic Value of Death

In this section, I shall explain why the Causal Consequence Version implies that you and I respectively act wrongly in Case One, only if we also assume that the state of affairs that Victim dies at t_1 is intrinsically bad. (It will also become evident why the Causal Consequence Version has to make the same assumption to handle Case Two.) I shall also suggest that this is an implausible suggestion.

Consider again Case One. I have already established that the state of affairs that Victim dies at t_1 is a causal consequence of my action and your action, respectively, according to Causation. But what would the causal consequences of our alternatives be? Let's focus on my action. Suppose that the only thing I could have done apart from shooting would be to stand still. (I also assume that the closest possible world to the actual world where I do not shoot Victim is the world where I stand still and you still shoot.) Moreover, suppose that Victim does not enjoy an ice cream at t_2 as a consequence of us shooting him. In that case, the state of affairs that Victim enjoys an ice cream at t_2 would obviously not be a causal consequence of you standing still, according to Causation since that state of affairs does not obtain in that possible world.

Consequently, the Causal Consequence Version implies that I do not act wrongly in Case One since the state of affairs that Victim dies at t_1 is neither good nor bad for him; and there is no good state of affairs associated with my alternative action (such as – for example – the state of affairs that Victim enjoys an ice cream at t_2). So, the Causal Consequence Version, does not handle cases like Case One any better than the Standard Version. Of course, if it were the case that the state of affairs that Victim dies at t_1 is intrinsically bad for him, the Causal Consequence Version would imply that you and I, respectively, act wrongly in Case One (and you alone in Case Two).

Is the state of affairs that Victim dies at t_1 intrinsically bad for him? Although it might initially seem to have some plausibility that this state of affairs is intrinsically bad for Victim (it concerns his death, after all), Hedonism and Preferentialism implies that it is not. According to Hedonism – i.e., the family of views that hold that a state of affairs is intrinsically bad in virtue of involving someone's being in pain – the state of affairs that Victim dies at t_1 is not intrinsically bad for Victim since it does not involve him being in pain. Of course, the state of affairs that Victim dies *in pain* at t_1 is intrinsically bad for Victim, according to Hedonism. But this state of affairs does not obtain in Case One. According to Preferentialism – i.e., the family of views that hold that a state of affairs is intrinsically bad in virtue of involving someone's preferences being frustrated – the state of affairs that Victim dies at t_1 is not intrinsically bad for Victim either since it does not involve Victim's preferences

being frustrated. In fact, almost everyone that thinks that death is bad for the person who dies, thinks that death is bad for the person only in virtue of depriving that person of future intrinsic value.³³ Now, although the fact that someone dies at a certain point implies that she is deprived of the intrinsic value she would otherwise have had beyond that point, this fact does not – as we have seen earlier in this paragraph – imply that it is itself intrinsically bad.

There are of course other views on well-being apart from Hedonism and Preferentialism; for example, Perfectionism – i.e., the family of views that hold that a state of affairs is intrinsically bad in virtue of involving someone's failing to flourish. It might be that the state of affairs that Victim dies at t_1 is intrinsically bad for him, according to some version of Perfectionism, or it might be that this state of affairs is bad for Victim, according to some other view on well-being I have not mentioned. However, it would be better not to rely on these views since they are quite implausible, and since it is not clear that they imply that state of affairs that Victim dies at t_1 intrinsically bad for him.

Finally, even if it were the case that Victim's death is intrinsically bad for him or just intrinsically bad, the Causal Consequence Version would nevertheless not be very attractive. Consider again Case One but suppose this time that the only alternative I have, apart from shooting Victim, is to break the leg of Victim's sister. In that case, the Casual Consequence Version would imply that it is right to break the leg of Victim's sister since the disvalue of death arguably (that is, if you think death has negative intrinsic value) is higher than the disvalue of a broken leg. Intuitively, it is wrong to break the leg of Victim's sister, however. Since Victim would have been killed anyway, breaking the leg of Victim's sister simply adds misery to the world.

³³ David Benatar has expressed some sympathy in favour of the view that death is intrinsically bad for the person who dies in a talk he gave at a conference at Syracuse University. ("Deprived and Annihilated", Thursday, May 19, IAPDD Conference, 2016.) He suggested that annihilation might be intrinsically bad for the person who dies.

Chapter Five, The Non-Standard Version

1. Introduction

In the previous chapters I have often noted that the Standard Version of Act-Consequentialism implies that neither you nor I act wrongly in Case One and Case Two (i.e., the cases where you and I shoot and kill a third person called “Victim”). Moreover, I have argued that both strategies that the proponent of the Standard Version might employ to handle Case One and Case Two are problematic (i.e., accepting that that the implications of Standard Version in Case One and Case Two are counterintuitive but explain away these intuitions, and denying that these implications are counterintuitive, respectively). In the previous chapter, I wondered whether the Causal Consequences Version of Act-Consequentialism might be more successful than the Standard Version but concluded that it has severe shortcomings. For example, the Causal Consequences Version implies that you and I respectively act wrongly in Case One and that you alone act wrongly in Case Two, only if we implausibly assume that the state of affairs that Victim dies at t_1 is intrinsically bad.

In this chapter I am going to defend another version of Act-Consequentialism that I call “the Non-Standard Version”. This theory does not claim, as the Standard Version of Consequentialism does, that an action is wrong only in virtue of having an alternative whose outcome is intrinsically better. Rather, it claims that an action is wrong in virtue of having an alternative whose outcome is intrinsically better or being a redundant negative difference maker. The Non-Standard Version does have the intuitively correct implications in cases like Case One and Case Two. But apart from that it has the same implications as the Standard Version. There are some serious objections to the Non-Standard Version, however. I discuss some in this chapter and some in the next. As I said in the introduction to this thesis, my aim is to show that the Non-Standard Version is plausible enough to be worthy of consideration.

This chapter is organised as follows: In section 2, I introduce the Non-Standard Version, or rather its wrongness criterion, and show that it has the intuitively correct implications in a number of important cases. I also discuss a few worries here. In section 3, I discuss the rest of the Non-Standard Version. Here I also consider some further objections. Finally, in section 4, I discuss a principle I call the “Parfittian Principle”. This principle might be thought to be an attractive alternative to the Non-Standard Principle. I shall argue that the Non-standard Version is a bit more attractive.

Before I get going, I want to warn the reader. I am going to discuss quite a few cases in this chapter and the next. Unfortunately, this discussion will at points be somewhat repetitive. There is no way around that, however. I am claiming that the Non-Standard Version has the intuitively correct implications in Case One and Case Two. There would be little cause for celebration if there were other cases – similar to Case One and Case Two – where the Non-Standard Version does not have the intuitively correct implications.

2. The Non-Standard Version

In this section, I shall present the Non-Standard Version of Act-Consequentialism, or more specifically, I shall introduce its criterion of wrongness. As I said in the introduction to this chapter, I shall discuss the rest of the Non-Standard Version in section 3. In this section I shall also demonstrate that the Non-Standard Version has the intuitively correct implication in a number of cases, some of which we have discussed in previous chapters and some of which we have not yet considered.

This section is arranged as follows: In section 2.1, I introduce the Non-Standard Version (its criterion of wrongness) and show that it has the intuitively correct implications in Case One. In section 2.2, I show that the Non-Standard Version has the intuitively correct implications in Case Three and in Case Four. (In the previous chapters I have discussed a number of other cases. What about them? Later in this chapter, I discuss Case Five. The remaining ones, Case Two and Case Six, I discuss in the coming chapter.) In section 2.3, I show that the Non-Standard Version avoids moral dilemmas. In section 2.4, I show that the Non-Standard Version does not attribute wrongness where it is not due. Finally, in section 2.5, I discuss two worries concerning the Non-Standard Version.

2.1 The Non-Standard Version and Case One

In this section I shall introduce the Non-Standard Version's criterion of wrongness and I shall also show that it has the intuitively correct implications in Case One.

Here is the Non-Standard Version's criterion of wrongness:

The Non-Standard Version (its criterion of wrongness)

An action *a* is wrong, if and only if,

(1) *a* has an alternative whose outcome would be intrinsically better than that of *a*, or

(2)(a) *a* is a redundant negative difference maker, and

(b) *a* has an alternative that would not be a redundant negative difference maker

I have already explained important parts of this principle. In Chapter Two, I introduced and interpreted the Standard Version of Act-Consequentialism. Clause (1) of this principle is identical with the Standard Version's criterion of wrongness. In Chapter Two, I also stipulated what redundant negative difference makers are. But let me present that stipulation again.

Here it is:

Redundant Negative Difference Making

An action *a* is a redundant negative difference maker, if and only if, there is a combination of actions *C* such that

- (1) there is no alternative to *a* such that the outcome would be intrinsically better if that alternative and the combination *C* had been performed, and there is at least one alternative to *a* such that the outcome would not be intrinsically worse if that alternative and the combination *C* had been performed,
- (2) there is no alternative to the combination *C* such that the outcome would be intrinsically better if that alternative and *a* had been performed, and there is at least one alternative to the combination *C* such that the outcome would not be intrinsically worse if that alternative and *a* had been performed,
- (3) there is an alternative to the combination *a* and *C* such that the outcome would be intrinsically better if that alternative had been performed.

As I said in Chapter Two (section 2.3), a combination of actions may either be a compound action or a plurality of actions. Furthermore, the alternatives to a combination of actions are a function of its simple actions. For example, provided that both you and I, respectively, could either shoot or abstain from shooting, the alternatives to the combination "my shooting Victim, and your shooting Victim" are the following three: "my shooting Victim, and your abstaining from shooting Victim", "my abstaining from shooting Victim, and your shooting Victim", and "my abstaining from shooting Victim, and your abstaining from shooting Victim."

I have not yet said anything about clause (2)(b) of the Non-Standard Version. I shall explain why it is included in section 2.3, but let us focus on Case One for now. Here it is:

Case One

You and I shoot and kill Victim. Our bullets arrive at the same time and each shot would have killed Victim by itself. Victim would not have been killed, if neither of us had pulled the trigger. Apart from shooting Victim, you and I have only one alternative each: each could also have abstained

from shooting Victim. Furthermore, each of us would have performed his or her alternative – i.e., each of us would have abstained from shooting Victim – if he or she had not shot him. We act independently of one another, Victim’s level of well-being is not affected by either of the shots taken by itself (e.g., Victim does not experience any pain), and no one else is affected. Finally, Victim would have gone on to live a life well worth living, if neither you nor I had shot him.

Neither my action nor your action satisfies clause (1) of the Non-Standard Version: If I had not shot Victim you would still have done so, and the other way around. (Of course, we already knew that since clause (1) of the Non-Standard Version is identical with the Standard Version’s criterion of wrongness.) So, if the Non-Standard Version is able to handle Case One it must be because my action and your action, respectively, in Case One, satisfy clause (2) of the Non-Standard Version.

Are both my action and your action, respectively, redundant negative difference makers in Case One, according to Redundant Negative Difference Making? Yes, consider your action, for instance. There is a combination of actions – namely, my action – such that the three clauses of Redundant Negative Difference Making are satisfied. Clause (1) is satisfied since you only have one alternative – i.e., your abstaining from shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse if you had performed that alternative and I had shot him since (again) one shot is sufficient to kill him. Moreover, clause (2) is satisfied since I only have one alternative – i.e., my abstaining from shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse if I had performed that alternative and you had shot Victim since one shot is sufficient to kill him. Finally, clause (3) is satisfied since we have an alternative – i.e., the combination “my abstaining from shooting Victim, and your abstaining from shooting Victim” – such that the outcome would have been intrinsically better, if we had performed that combination: Victim would under those circumstances not have been killed.

A parallel argument establishes that my action is a redundant negative difference maker, according to Redundant Negative Difference Making. (Of course, we already knew that my action and your action, respectively, are redundant negative difference makers in Case One, according to Redundant Negative Difference Making, since I also demonstrated this in Chapter Two, section 3.1.)

But do both you and I, respectively, have an alternative that would not be a redundant negative difference maker, according to Redundant Negative Difference Making? Yes, consider, for example, your alternative (i.e., your abstaining from shooting Victim). Consider the closest possible world where you perform that alternative. Clause (2) is not satisfied in that world since the outcome would be intrinsically better if I had abstained from shooting Victim and

you had abstained from doing so. Parallel remarks show that my alternative is not a negative redundant difference maker, according to Redundant Negative Difference Making. Consequently, the Non-Standard Version implies that you and I respectively act wrongly in Case One.

2.2 Case Three and Case Four

So, the Non-Standard Version has the intuitively correct implications in Case One. What about Case Three and Case Four, i.e., the cases that are different from Case One in that there are more than two actions involved? The Non-Standard Version has the intuitively correct implications in these two cases.

Consider first Case Three:

Case Three

You, Gertrude, and I shoot and kill Victim. Our bullets arrive at the same time and each shot would have killed Victim by itself. Victim would not have been killed, if none of us had pulled the trigger. Apart from shooting Victim, each of us has only one alternative: each could also have abstained from shooting Victim. Furthermore, each of us would have performed his or her alternative – i.e., each of us would have abstained from shooting Victim – if he or she had not shot him. We act independently of one another, Victim’s level of well-being is not affected by either of the shots taken by itself (e.g., Victim does not experience any pain), and no one else is affected. Finally, Victim would have gone on to live a life well worth living, if none of us had shot him.

In Chapter Two I noted that none of us acts wrongly in Case Three, according to the Standard Version: None of us has an alternative whose outcome is intrinsically better. But we do seem to act wrongly. (Case Three is clearly relevantly similar to Case One in that respect.) However, as I demonstrated in the same chapter, our actions in this case are redundant negative difference makers, according to Redundant Negative Difference Making (Chapter Two, section 3.2).

But is clause (2)(b) of the Non-Standard Version satisfied? Yes! Gertrude, you and I have an alternative that would not be a redundant negative difference maker, according to Redundant Negative Difference Making. Consider Gertrude’s alternative (i.e., Gertrude’s abstaining from shooting Victim), for example. Go to the closest possible world where Gertrude performs that alternative. Clause (2) is not satisfied in that world since the outcome would be intrinsically better if both you and I had abstained from shooting Victim and Gertrude had done so too.

Parallel remarks show that my alternative and your alternative would not be redundant negative difference makers either, according to Redundant Negative Difference Making. So, the Non-Standard Version implies that Gertrude, you and I respectively act wrongly in Case Three.

Let us now turn to Case Four:

Case Four

You, Gertrude and I together shoot and kill Victim. Our bullets arrive at the same time but we are dealing with a very odd victim: he won't die unless two bullets hit him simultaneously (but two bullets would do the job just fine). Victim would not have been killed, if none of us had pulled the trigger. Apart from shooting Victim, each of us has only one alternative: each could also have abstained from shooting Victim. Furthermore, each of us would have performed his or her alternative – i.e., each of us would have abstained from shooting Victim – if he or she had not shot him. We act independently of one another, Victim's level of well-being is not affected by either of the shots taken by itself (e.g., Victim does not experience any pain), and no one else is affected. Finally, Victim would have gone on to live a life well worth living, if none of us had shot him.

I also noted, in Chapter Two, that none of us acts wrongly in Case Four, according to the Standard Version: neither you, nor I, nor Gertrude has an alternative whose outcome would be intrinsically better. However, it seems intuitively clear in this case too, that we do act wrongly. But as I showed, in Chapter Two, our actions in Case Four are redundant negative difference makers (Chapter Two, section 3.2).

But do Gertrude, you and I have an alternative that would not be a redundant negative difference maker, according to Redundant Negative Difference Making? Yes, consider the alternative to your action, for instance. Go to the closest possible world where you perform that alternative. Clause (2) is not satisfied in that world since the outcome would be intrinsically better if, for example, I had abstained from shooting Victim (or if Gertrude and I had abstained from shooting Victim) and you had abstained from shooting him. Parallel remarks show that Gertrude and I respectively have an alternative that would not be a redundant negative difference maker either, according to Redundant Negative Difference Making.

2.3 Moral Dilemmas

Why does the Non-Standard Version's criterion of wrongness include clause (2)(b)? Well, there are cases where an action is a redundant negative difference maker, and where all its alternatives would be redundant negative differ-

ence makers, according to Redundant Negative Difference Making, and without clause (2)(b), the Non-Standard Version (or, strictly speaking, the version of the Non-Standard Version that doesn't include this clause) would in these cases imply moral dilemmas, i.e., that the agents involved act wrongly whatever they do. I want to avoid that.

Here is a case like that:

Case Seven

Case Seven is like Case One except that the only alternative I had was to shoot Victim's sister. If I had chosen to shoot Victim's sister, my bullet would have hit her at the same time as Gertrude's bullet (who is actually shooting Victim's sister) would have hit her. My shot would have been sufficient to kill Victim's sister and Gertrude's shot actually kills her. Apart from shooting Victim's sister, Gertrude could only have abstained from shooting Victim's sister and that is also what she would have done if she had not shot Victim's sister. Gertrude acts independently of you and me, and the other way around. Finally, no further person is affected by what Gertrude does. (Note that Case Seven is like Case One – except for the stated exceptions – and that I did assume that you and I only had one alternative, that we would perform that alternative if we had not shot Victim, and that we do not affect any further person.)³⁴

According to Redundant Negative Difference Making, my action is a redundant negative difference maker. There is a combination of actions – your action – such that the three clauses of Redundant Negative Difference Making are satisfied. Clause (1) is satisfied since I only have one alternative – i.e., my shooting Victim's sister – and the outcome would have been neither intrinsically better nor intrinsically worse if I had shot Victim's sister, and you had shot Victim. Both Victim and Victim's sister would under those circumstances still have been killed. Victim would have been killed by you, and Victim's sister would have been killed by Gertrude: One shot is enough to kill Victim and one shot is also enough to kill Victim's sister.

³⁴ Erik Carlson brought cases like Case Seven to my attention. The case he came up with is a lot neater than Case Seven. I decided to use Case Seven because it sort of grows naturally out of Case One. But let me in any event present Carlson's suggestion: There are three shooters. The optimal outcome is achieved, if and only if, only B or only A and C shoot. B and C do not shoot regardless of what A does. A's shooting, *a*, is then a redundant negative difference maker: the outcome would have been neither intrinsically better nor intrinsically worse if *-a* and *-b* had been performed, and the outcome would have been neither intrinsically better nor intrinsically worse if *b* and *a* had been performed, but there is a better alternative to *a* & *-b*, namely *-a* & *b*. A's not shooting, *-a*, is also a redundant negative difference maker: the outcome would have been neither intrinsically better nor intrinsically worse if *a* and *-c* had been performed, and the outcome would have been neither intrinsically better nor intrinsically worse if *c* and *-a* had been performed, but there is a better alternative to *-a* & *-c*, namely *a* & *c*.

Moreover, clause (2) is satisfied since you only have one alternative – i.e., your abstaining from shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse, if you had not shot Victim, and I had shot him. Both Victim and Victim’s sister would under these circumstances still have been killed. Victim would have been killed by me, and Victim’s sister would have been killed by Gertrude. Finally, clause (3) is satisfied since you and I have an alternative – i.e., the combination “my shooting Victim’s sister, and your abstaining from shooting Victim” such that the outcome would have been intrinsically better, if we had performed that combination: Although Victim’s sister would still have been killed by Gertrude and me, Victim would not have been killed since it takes at least one shot to kill him.

However, my alternative – i.e., my shooting Victim’s sister – is also a redundant negative difference maker, according to Redundant Negative Difference Making. Go to the closest possible world where I shoot Victim’s sister. There is a combination of actions – Gertrude’s action – such that the clauses of Redundant Negative Difference Making are satisfied here. Clause (1) is satisfied since I only have one alternative – i.e., my shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse, if I had shot Victim and Gertrude had shot Victim’s sister. Both Victim’s sister and Victim would still have been killed. Victim would still have been killed by both me and you (respectively), and Victim’s sister would still have been killed by Gertrude. Furthermore, clause (2) is satisfied since Gertrude only has one alternative – i.e., Gertrude’s abstaining from shooting Victim’s sister” – and the outcome would have been neither intrinsically better nor intrinsically worse, if Gertrude had not shot Victim’s sister and I had done so. Both Victim’s sister and Victim would still have been killed. Victim would have been killed by you, and Victim’s sister would have been killed by me. Finally, clause (3) is satisfied since Gertrude and I have an alternative – i.e., the combination “my shooting Victim, and Gertrude’s abstaining from shooting Victim’s sister” such that the outcome would have been intrinsically better, if we had performed that combination: Although Victim would still have been killed by you and me, Victim’s sister would not have been killed since it takes at least one shot to kill her.

So, my action is a redundant negative difference maker, and is therefore wrong, according to the version of the Non-Standard Version that does not include clause (2)(b). Furthermore, my alternative (I have only one) would be a redundant difference maker, and would consequently be wrong, according to the version of the Non-Standard Version that does not include clause (2)(b). Finally, since my action is wrong and my only alternative would be wrong, according to the version of the Non-Standard Version that does not include clause (2)(b), this version implies a moral dilemma, i.e., that I act would act wrongly whichever of my alternatives I would perform.

But why do I think it is important to avoid moral dilemmas? Personally, I am not convinced that there are no moral dilemmas. For example, if Gertrude has to choose between killing her son and her daughter, it does not seem implausible to say that she acts wrongly whatever she does. She just finds herself in one of those horribly tragic situations that people sometimes find themselves in. On the other hand, it does not seem implausible to say that if a person acts wrongly, then that person should perform another action, one that is not wrong.

Or perhaps not. We might perhaps say that although Gertrude would act wrongly whichever of her children she would choose to kill, she is not to be blamed for killing one of them (since it was unavoidable). However, this opens up a whole can of worms. What is the relationship between wrongness and blameworthiness? If these conflict, should we strive to avoid blame or should we strive to avoid acting wrongly? I am not going to pursue this issue here. Instead, I note that consequentialists do not normally recognise moral dilemmas; and since the Non-Standard Version is a version of consequentialism, it is better to keep that feature.

2.4 Innocent Agents

As we shall see when we discuss the Parfitian Principle in section 4, Parfit struggled with cases like this:

Case Eight

Case Eight is exactly as Case Four except that there is yet a further agent – Fred Astaire – who dances ferociously in the far distance. Fred Astaire has only one alternative: he could have abstained from dancing, and this is also what he would have done, if he had not danced. Furthermore, Fred Astaire acts independently of us and his action does not affect anyone at all.

Parfit did not himself present a principle that has the intuitively correct implications in this case, namely, that you and I and Gertrude, respectively, act wrongly, and that Fred Astaire does not. (The principle I call the “Parfitian Principle” does handle this case, though.) In this section I want to show that the Non-Standard Version has the intuitively correct implications in this case. In section 2.2, I demonstrated that the Non-Standard Version implies that you, I and Gertrude act wrongly. So, it remains to be shown that the Non-Standard Version implies that it is not the case that Fred Astaire acts wrongly.

Fred Astaire’s action does not satisfy clause (1) of the Non-Standard Version. The outcome would not have been intrinsically better, if he had acted differently. His action does not satisfy clause (2) either. It is not a negative redundant difference maker.

Suppose first that the variables *C* pick out one of our actions, for example, yours. (The variable *a* picks out Fred Astaire's action.) In that case, clause (3) of Redundant Negative Difference Making is not satisfied. You and Fred Astaire do not have an alternative whose outcome would be intrinsically better. For example, the outcome would not have been intrinsically better if neither you nor Fred Astaire had acted as you did since Victim would under those circumstances still have been killed by me and Gertrude. Suppose now that the variable *C* instead pick out several of our actions, for example, mine and yours. (Again, the variable *a* picks out Fred Astaire's action.) In that case, clause (2) of Redundant Negative Difference Making is not satisfied. There is an alternative to what you and I together do whose outcome would be intrinsically better: The outcome would have been intrinsically better if neither you nor I had shot Victim and Fred Astaire still had danced since Victim would not have been killed under those circumstances: Gertrude's shot would not be sufficient to kill Victim since it takes two shots to kill him here.

2.5 Two Worries

In this section I shall briefly discuss two worries that proponents of the Standard Version might have with respect to the Non-Standard Version.

First, they might worry that I have added a foreign element to consequentialism by the inclusion of clause (2). I think that there is a close affinity between the two clauses, however. Clause (1) of the Non-Standard Version is concerned with whether there is an alternative to a particular action whose outcome is intrinsically better than those of the action in question and clause (2) of the Non-Standard Version is crucially concerned with whether there is an alternative to a combination of particular actions whose outcome is better than those of the combination in question. Both these concerns are consequentialist at heart.

Second, they might worry that the Non-Standard Version is too complex to be an attractive alternative to the Standard Version. (I haven't yet introduced the other parts of the Non-Standard Version!) They might concede that the fact that the Non-Standard Version does seem to have the intuitively correct implications in cases like Case One is some evidence in favour of the Non-Standard Version, but point out that there is as weighty evidence that speaks against the Non-Standard Version and that it therefore – on balance – seems better to keep the simpler view. As we shall see, there are quite many case that where it is not so clear that the Non-Standard Version has the intuitively correct implications.

I do not have a knock-down reply to this worry. I shall simply try to show in what follows that the objections to the Non-Standard Version that we will encounter in this chapter and the next are manageable. Inevitably, this will at some points – when discussing a troubling case – be no more than suggesting

that the implications of the Non-Standard Version in the troubling case at hand aren't counterintuitive.

3. Other Deontic Categories

In this section I briefly shall introduce the Non-Standard Version's criterion of rightness and its criterion of obligatoriness. Moreover, I shall discuss some objections to the Non-Standard Version. For instance, the Non-Standard Version does not, as we shall see, say that it is obligatory to perform an action that is what I call a "redundant positive difference maker" (i.e., an action that would be a positive difference maker if it were not for other actions involved). But as we have seen, the Non-Standard Version claims that it is wrong to perform an action that is a redundant negative difference maker. But how can this asymmetry between redundant positive difference makers and redundant negative difference makers be justified? I shall also discuss some other difficulties.

I have arranged this section in the following manner: In section 3.1, I introduce the Non-Standard Version's criterion of rightness and its criterion of obligatoriness. In section 3.2, I say what a redundant positive difference maker is. Here I also present a case that involves redundant positive difference makers. In section 3.3 I explain why I think the asymmetry between redundant positive difference makers and redundant negative difference makers may be justified. In section 3.4, I discuss a case where an action is both a redundant negative difference maker and a redundant positive difference maker. Finally, in section 3.5, I discuss a case where an action that is a redundant negative difference maker has as its only alternative a redundant positive difference maker.

3.1 Rightness and Obligatoriness

In this section, I introduce the Non-Standard Version's criterion of rightness and obligatoriness. To make these criteria easier to appreciate, I use the Non-Standard Version's account of wrongness to describe them.

Here is the Non-Standard Version's criterion of rightness:

The Non-Standard Version (its criterion of rightness)

An action *a* is *right*, if and only if, it is not wrong, according to the Non-Standard Version's criterion of wrongness.

And here is its criterion of obligatoriness:

The Non-Standard Version (its criterion of obligatoriness)

An action *a* is *obligatory*, if and only if, each alternative to *a* would be wrong, according to the Non-Standard Version's criterion of wrongness.

3.2 Redundant Positive Difference Making

I shall characterise Redundant Positive Difference Making as follows:

Redundant Positive Difference Making

An action *a* is a positive difference maker, if and only if, there is a combination of actions *C* such that

- (1) there is no alternative to *a* such that the outcome would be intrinsically better if that alternative and the combination *C* had been performed, and there is at least one alternative to *a* such that the outcome would not be intrinsically worse if that alternative and the combination *C* had been performed,
- (2) there is no alternative to the combination *C* such that the outcome would be intrinsically better if that alternative and *a* had been performed, and there is at least one alternative to the combination *C* such that the outcome would not be intrinsically worse if that alternative and *a* had been performed,
- (3) there is an alternative to the combination *a* and *C* such that the outcome would be intrinsically *worse* if that alternative had been performed.

Redundant Positive Difference Making is similar to Redundant Negative Difference Making. However, there is of course a crucial difference. Clause (3) of Redundant Positive Difference Making requires that there is an alternative whose consequences are intrinsically *worse*, whereas clause (3) of Redundant Negative Difference Making requires that there is an alternative whose consequences are intrinsically *better*.

Let us consider a case where the actions involved are redundant positive difference makers:

Case Nine

A poisonous snake has bitten Victim. If Victim is not given one dosage of snake anti-venom serum within one hour, he will die. You and I have one dosage each of the required kind (i.e., each dosage is sufficient to save Victim's life). Before one hour has passed, you and I, respectively, inject Victim with one dosage each. Everything goes well, Victim survives. Apart from injecting Victim, you and I have only one alternative each: each could

also have abstained from injecting Victim. Furthermore, each of us would have performed his or her alternative – i.e., each of us would have abstained from injecting Victim – if he or she had not injected him. We act independently of one another, Victim’s level of well-being is not affected by either of the injections taken by itself (e.g., Victim does not experience any pain), and no one else is affected. Finally, Victim’s continued life is a life well worth living.

Now, consider your action. There is a combination – namely, my action – such that the clauses of Redundant Positive Difference Making are satisfied. Clause (1) is satisfied since you only have one alternative – i.e., your abstaining from injecting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse if you had abstained from injecting Victim and I had injected him since one dosage is sufficient to save him. Furthermore, clause (2) is satisfied since I only have one alternative – i.e., my abstaining from injecting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse if I had abstained from injecting Victim and you had injected him since (once more) one dosage is sufficient to save him. Finally, clause (3) is satisfied since we have an alternative – i.e., the combination “my abstaining from injecting Victim, and your abstaining from injecting Victim” – such that the outcome would have been intrinsically worse, if we had performed that combination: Victim would under those circumstances not have been saved. Consequently, your action is a redundant positive difference maker, according to Redundant Positive Difference Making, and a parallel argument establishes that my action is one too.

3.3 The Symmetry Objection

The symmetry objection essentially complains that if it is wrong to perform an action that is a redundant negative difference maker, it ought to be obligatory to perform an action that is a redundant positive difference maker, and hence wrong to fail to perform an action that is a redundant positive difference maker.

Consider this case:

Case Ten

This case is exactly like Case Nine except that you do not inject Victim with the serum. You just stare at us. (However, since I inject Victim with the required dose in time, he survives.)

Now, note that you don’t act wrongly, according to the Non-Standard Version. You do not have an alternative whose consequences are intrinsically better. What I do is enough to save Victim. Moreover, your action is obviously not a

negative redundant difference maker. (Clause (3) of Redundant Negative Difference Making is not satisfied since there is not an alternative to what you and I do such that the outcome would be intrinsically better if that alternative had been performed: I am independently injecting Victim with the required dose.) However, should we not say that you do act wrongly here? If it is wrong to perform actions that are negative redundant difference makers, why – we may ask – is it not obligatory to perform actions that are *positive* redundant difference makers (or wrong to fail to perform them)? There seem to be a curious asymmetry here.

I agree that there is a somewhat curious asymmetry here. However, I shall nevertheless leave things as they are. Although it is possible to accommodate the view that it is obligatory to perform actions that are positive redundant difference makers, I do not think that it is worth the trouble. It would require a principle that is a lot more complicated than the Non-Standard Version, and I am not convinced that it is counterintuitive that you do not act wrongly here.

3.4 A Mixed Case

As I said, there are cases where an action is both a redundant negative difference maker and a redundant positive difference maker.

Here is a case like that:

Case Eleven

You and I shoot and kill Victim. Our bullets arrive at the same time and each shot would have killed Victim by itself. Victim wouldn't have been killed, if neither of us had shot him. Furthermore, my shot together with Gertrude's shot – who is out hunting – scare off a wild animal (whose welfare level is not affected) that would have killed Victim's sister, if our shooting had not disturbed it. Each shot – i.e., my shot and Gertrude's shot, respectively – would have scared off the animal by itself. Each of us acts independently and each of us had one alternative apart from shooting, i.e., abstaining from shooting, which we would have performed, if we had not fired our guns. No one else is affected, and Victim would have gone on to live a life well worth living, if you and I had not attacked him. Victim's sister would have been deprived of her life if neither I nor Gertrude had acted as we did, and her continued life is well worth living.

Note first that my action is a redundant negative difference maker, according to Redundant Negative Difference Making. There is a combination of actions – namely, your action – such that the three clauses of Redundant Negative Difference Making are satisfied. Clause (1) is satisfied since I only have one alternative – i.e., my abstaining from shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse if I had

abstained from shooting Victim and you had shot him since one shot is sufficient to kill him and Gertrude's shot is sufficient to scare of the wild animal. Furthermore, clause (2) is satisfied since you only have one alternative – i.e., your abstaining from shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse if you had abstained from shooting Victim and I had shot him since one shot is sufficient to kill Victim and since my shot and Gertrude's shot are sufficient to scare of the wild animal. Finally, clause (3) is satisfied since we have an alternative – i.e., the combination “my abstaining from shooting Victim, and your abstaining from shooting Victim” – such that the outcome would have been intrinsically better, if we had performed that combination: Victim would under those circumstance not have been shot, and Gertrude's shot would still have scared off the wild animal.

Furthermore, my action is also a redundant positive difference maker, according to Redundant Positive Difference Making. There is a combination – namely, Gertrude's action – such that the clauses of Redundant Positive Difference Making are satisfied. Clause (1) is satisfied since I only have one alternative – i.e., my abstaining from shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse if I had abstained from shooting Victim and Gertrude had fired her gun: You would still have killed Victim and Gertrude would still have scared the animal. Moreover, clause (2) is satisfied since Gertrude only has one alternative – i.e., Gertrude's abstaining from shooting – and the outcome would have been neither intrinsically better nor intrinsically worse if Gertrude had abstained from firing her gun and I had shot Victim: You and I would still have killed Victim and I would still have scared of the wild animal. Finally, clause (3) is satisfied since Gertrude and I have an alternative – i.e., the combination “my abstaining from shooting Victim, and Gertrude's abstaining from shooting” – such that the outcome would have been intrinsically worse, if we had performed that combination: The wild animal would then have killed Victim's sister.

Now, the Non-Standard Version implies that I act wrongly. My action is a redundant negative difference maker that also has an alternative that is not a redundant negative difference maker. But is that really what we want to say?

I shall not revise the Non-Standard Version, although it would not be very difficult to do so. For example, to revise the Non-Standard Version's criterion of wrongness one would just have to add the following clause of the Non-Standard Version:

(2)(c) *a* is not also a redundant positive difference maker

However, I do not think that it would be worthwhile to do so. Again, it is not clear to me that this is counterintuitive. I therefore think that it is better to leave things as they are.

3.5 A Further Case

In section 2.3 I discussed a case (i.e., Case Seven) where an action and its alternative were redundant negative difference makers, according to Redundant Negative Difference Making. In this section I am going to discuss a case where an action is a redundant negative difference maker, according to Redundant Negative Difference Making, but where its alternative would be a redundant positive difference maker, according to Redundant Positive Difference Making.

Here it is:

Case Twelve

Case Twelve is like Case One except that if I had abstained from shooting Victim I would have shot a rabid dog instead (whose welfare level is not affected) that was just about to bite Victim's sister. If I had chosen to shoot the rabid dog, my bullet would have hit the dog at the same time as Gertrude's bullet (who is actually shooting the dog) would have hit her. My shot would have been sufficient to kill the dog and Gertrude's shot actually kills her. If the dog had bitten Victim's sister, she would have died in terrible pain. Gertrude acts independently of you and me, and the other way around. Finally, no further person is affected by what Gertrude does, and Gertrude had one alternative: she could also have abstained from shooting the dog, and this is also what she would have done, if she had not shot the dog. (Note that Case Twelve is like Case One – except for the stated exceptions – and that I did assume that neither you nor I affect any further person there and that you and I also had one alternative: we could have abstained from shooting Victim, and that is also what we would have done, if we had not shot him.)

My action is a redundant negative difference maker, according to Redundant Negative Difference Making. There is a combination – your action – such that the clauses of Redundant Negative Difference Making are satisfied. Clause (1) is satisfied since I only have one alternative – i.e., my shooting the rabid dog – and the outcome would have been neither intrinsically better nor worse, if I had not shot Victim and instead shot the rabid dog, and you had shot Victim: Victim would still have been killed and the dog would still have been prevented from biting Victim's sister. Furthermore, clause (2) is satisfied since you only have one alternative – i.e., your abstaining from shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse, if you had not shot Victim, and I had shot him but not the dog: Victim would also under these circumstances still have been killed and the dog would also under these circumstances still have been prevented from biting Victim's sister since I would have shot Victim and Gertrude would have shot the dog.

Finally, clause (3) is satisfied since you and I have an alternative – i.e., the combination “my shooting the rabid dog, and your abstaining from shooting Victim” – such that the outcome would have been intrinsically better, if we had performed that combination: Although the dog would still have been prevented from biting Victim’s sister, Victim would not have been killed since it takes at least one shot to kill him.

However, my alternative – i.e. my shooting the rabid dog – would be a redundant positive difference maker, according to Redundant Positive Difference Making. Consider the closest possible world where I shoot the rabid dog. There is a combination – Gertrude’s action – such that the clauses of Redundant Positive Difference Making are satisfied. Clause (1) is satisfied there since I only have one alternative – my shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse, if I had not shot the dog and Gertrude had shot her. Victim would still have been killed and the dog would still have been prevented from biting Victim’s sister. Victim would have been killed by you and me, and the rabid dog by Gertrude. Furthermore, clause (2) is satisfied since Gertrude only has one alternative – i.e., Gertrude’s abstaining from shooting the dog – and the outcome would have been neither intrinsically better nor intrinsically worse, if Gertrude had not shot the rabid dog and I had done so. Victim would still have been killed and the dog would still have been prevented from biting Victim’s sister. Victim would have been killed by you, and the rabid dog by me. Finally, clause (3) is satisfied since Gertrude and I have an alternative such that the outcome would have been intrinsically worse if we had performed that alternative: the outcome would have been intrinsically worse if neither I nor Gertrude had shot the dog. Although Victim would still have been killed, Victim’s sister would then have been bitten by the rabid dog.

What does the Non-Standard Version imply in Case Twelve? Well, it implies that I act wrongly. My action would be a redundant negative difference maker that also has an alternative that would not be a redundant negative difference maker: My alternative would be a redundant positive difference maker. Is this what we want to say in this case? Well, I think we may say so. The reason I was “absolved” in Case Nine was that I had no alternative but to do an equally bad thing. Here I could instead do a good thing.

4. The Parfittian Principle

I turn now to the principle I call the “Parfittian Principle.” (I call this principle the “Parfittian Principle” because it is very similar to a principle proposed by

Parfit; but not similar enough – I think – to deserve being called “Parfit’s principle.”)³⁵ This principle is designed to handle cases like Case One, and it might be thought that it is an attractive alternative to the Non-Standard Version. It has also received a fair amount of attention recently (well, Parfit’s principle has). However, although the Parfittian Principle does have the intuitively correct implications in Case One and other similar cases (as I shall demonstrate further down), I shall argue that the Non-Standard Version nevertheless is preferable to the Parfittian Principle. There is one kind of case where the Parfittian principle does not have the desired implications.

In section 4.1, I introduce the Parfittian Principle. In this section, I also show that it has the intuitively correct implications in Case One. In section 4.2, I show that this principle has the intuitively correct implications in Case Three and Case Four. In section 4.3, I show that the Parfittian Principle has the intuitively correct implications in Case Seven and Case Eight. Finally, in section 4.4, I show that the Non-Standard Version is preferable to the Parfittian Principle. I present a case where the Parfittian principle does not have the intuitively correct implications, but where the Non-Standard Version has.

4.1 Introducing the Parfittian Principle, and Case One

In this section I introduce the Parfittian Principle, or rather its criterion of wrongness. (The point I want to make does not require a discussion of its criterion of rightness or its criterion of obligatoriness. I shall therefore ignore these two latter criteria.) Furthermore, I show that it has the correct implications in Case One.

Here it is:

The Parfittian Principle (its criterion of wrongness)

An action *a* is wrong, if and only if, either

- (1) there is an alternative to *a* whose outcome is intrinsically better than that of *a*, or
- (2) *a* is a member of a set of actions *S* such that there is an alternative to *S* whose outcome is intrinsically better than that of *S*, and
- (3) there is no proper subset of *S*, *P*, such that there is an alternative to *P* whose outcome is intrinsically better than that of *P*.

There are a few things I would like to note. First, what I just presented is not the complete principle. There is a fourth clause – clause (4) – that is needed in order to handle Case Seven. I shall introduce clause (4) when I discuss Case Seven, which I do at the very end of this section. (Clause (4) is a conjunct of

³⁵ For Parfit’s actual principle, see Derek Parfit, *Reasons and Persons*, Oxford, Clarendon Press (1984), 67–72, 82–83.

the second disjunct of the Parfittian Principle.) Second, I wrote that “*a* is a member of a set of actions *S* such that there is an alternative to *S* whose outcome is intrinsically better than that of *S*”. However, sets do not really have consequences. A more correct formulation would be “*a* is a member of a set of actions *S* such that there is an alternative to *S*, *T*, such that the outcome would have been intrinsically better had the members of *T* been performed instead of the members of *S*.” Similar remarks apply to the clause (3). But I leave things as they are for the sake of simplicity (relying on this remark to forestall confusion). Third, the alternatives to a set of actions mirror the alternatives to combination of actions discussed above. For example, the alternatives to the set of actions {my shooting Victim, and your abstaining from shooting Victim}, provided that both you and I, respectively, could either shoot or abstain from shooting, are the following: {my shooting Victim, and your abstaining from shooting Victim}, {my abstaining from shooting Victim, and your shooting Victim}, and {my abstaining from shooting Victim, and your abstaining from shooting Victim}.

As I said, the Parfittian Principle is different from the one Parfit suggests. Let me briefly mention some of the differences. First, Parfit’s principle is concerned with harm, and provides sufficient conditions. Second, Parfit’s principle involves the condition “*S* is the smallest set of which this is true” rather than as this principle does the criterion “there is no proper subset of *S*, *P*, such that there is an alternative to *P* whose outcome is intrinsically better than that of *P*” (which if translated into Parfit’s vernacular reads “there is no smaller set than *S* of which this is true”). Third, Parfit’s principle does not contain a clause corresponding to clause (4) of this principle (i.e., the clause I am going to introduce at the end of this section when I discuss Case Seven).

Consider now Case One, i.e., the case where you and I shoot and kill Victim simultaneously. Our actions are wrong in this case, according to the Parfittian Principle. Consider my action, for example. It is a member of a set of actions that has an alternative whose outcome is intrinsically better; namely the set containing my action and your action. The outcome would have been intrinsically better if neither you nor I had shot Victim. Furthermore, there is no proper subset of this set whose outcome is intrinsically better. Consider for example the set containing my action. The outcome would not have been intrinsically better if I had not shot Victim since you would still have shot him. My action is consequently wrong, according to the Parfittian Principle.

A parallel argument establishes that your action is wrong, according to the Parfittian Principle.

4.2 Case Three and Case Four

In this section I shall show that the Parfittian Principle has the intuitively correct implications in Case Three and Case Four. I discuss them in order.

Consider Case Three again, i.e., the case where you, Gertrude, and I shoot Victim (and where it takes one bullet to kill Victim). Our actions are wrong in Case Three, according to the Parfittian Principle. Consider for example Gertrude's action. It is a member of a set of actions that has an alternative whose outcome is intrinsically better; namely the set containing my action, your action and Gertrude's action. Moreover, there is no proper subset of that set that has an alternative whose outcome is intrinsically better. Consider the set containing Gertrude's action and my action, for instance. That set does not have an alternative whose outcome is intrinsically better since you would still have shot Victim had neither Gertrude nor I done so. It is evident that this also holds for the other proper subsets. Gertrude's action is consequently wrong, according to the Parfittian Principle.

Analogous arguments establish that your action and my action, respectively, are wrong according to the Parfittian principle.

Finally, your action, my action, and Gertrude's action are also wrong in Case Four, according to the Parfittian Principle, i.e., the case where you, Gertrude, and I shoot Victim but where it takes two bullets to kill Victim. Consider for example your action. It is a member of several sets of actions that have alternatives whose outcomes that are intrinsically better. It is, for example, a member of the set containing your action and Gertrude's action. The outcome would have been intrinsically better if you and Gertrude had acted differently since my shot is not sufficient by itself to kill Victim. Furthermore, there is no proper subset of that set – i.e., the set containing your action and Gertrude's action – that has an alternative whose outcome is intrinsically better. Victim would still have been killed, if, for example, only Gertrude had acted differently since both you and I would under those circumstances still have shot him. So you act wrongly in Case Four, according to the Parfittian Principle.

Analogous arguments establish that Gertrude's action and my action, respectively, are wrong according to the Parfittian principle.

There are two things it might be helpful to note about Case Four. First, the set containing your action, my action, and Gertrude's action does have an alternative whose outcome is intrinsically better. The outcome would have been intrinsically better if none of us had shot Victim. But there is a proper subset to this set that has an alternative whose outcome is intrinsically better; for example, the set discussed in the previous paragraph – i.e., the one containing your action and Gertrude's action. So my action does not come out wrong for that reason. Second, Parfit's own formulation of clause (3), i.e., "S is the smallest set of which this is true" would not render any of our actions wrong in Case Four. There is no set in the circumstances of which this is true: There are three equally small sets: {my shooting Victim, your shooting Victim}, {my shooting Victim, Gertrude's shooting Victim}, and {Gertrude's shooting Victim, your shooting Victim}. This is why I formulated clause (3) of the Parfittian Principle as I did.

4.3 Case Seven and Case Eight

I turn now to Case Seven and Case Eight. I shall show that the Parfittian Principle has the correct implication in these two cases. I discuss them in reversed order.

Consider Case Eight, i.e., the case where everything is exactly as in Case Four except for Fred Astaire (who dances ferociously in the far distance). Our actions (i.e., your action, my action and Gertrude's action) but not Fred Astaire's action, are wrong in this case, according to the Parfittian Principle. I have already shown that your action, my action, and Gertrude's action are wrong in Case Ten, according to the Parfittian Principle. I showed that when I showed that our actions in Case Four are wrong, according to that same principle. It remains to be shown the Parfittian Principle does not imply that Fred Astaire's action is wrong.

Fred Astaire's action is obviously a member of a set of actions, which has an alternative whose outcome is intrinsically better; namely the set containing your action, my action, Gertrude's action, and Fred Astaire's action. The outcome would for example have been intrinsically better if neither you nor I nor Gertrude had shot Victim, and Fred Astaire had continued dancing. However, there is also a proper subset to that set that has an alternative whose outcome is intrinsically better; namely the set containing your action, my action, and Gertrude's action. The outcome would have been intrinsically better if neither you nor I nor Gertrude had shot Victim. Fred Astaire's action is not therefore wrong, according to the Parfittian Principle.

Finally, I turn to Case Seven, i.e., the case where you and I shoot Victim, and where I could only have shot Victim's sister together with Gertrude if I had not shot Victim. I have not yet presented clause (4) of the Parfittian Principle. But without this clause – the Parfittian Principle – or rather the resulting principle, which we may call the “Parfittian Principle without clause (4)”, would imply a moral dilemma. My action is wrong, according to the Parfittian Principle without clause (4) since it is a member of a set of acts that has an alternative whose outcome is intrinsically better, and since there is no proper subset of which this is true. Moreover, these things are true of the only alternative I have (i.e., shooting Victim's sister). However, the following clause is a part of the Parfittian Principle (as I construe it):

- (4) there is an alternative to a that is not a member of a set of actions M such that
 - (a) there is an alternative to M whose outcome is intrinsically better than that of M , *and*
 - (b) there is no proper subset of M , Q , such that there is an alternative to Q whose outcome is intrinsically better than that of Q .

With this addition, the Parfittian Principle has the intuitively correct implications in Case Seven. My action (i.e., my shooting Victim) is not wrong, according to the Parfittian Principle. Consider my alternative, i.e., my shooting Victim's sister. It is a member of a set that is distinct from the set {my shooting Victim, your shooting Victim}: namely, the set {my shooting Victim's sister; Gertrude's shooting Victim's sister}. This latter set has an alternative whose outcome is intrinsically better. It is better if neither Gertrude nor I shoot Victim's sister: Victim's sister wouldn't under those circumstances get killed.

Finally, there is no proper subset of this set that has a better outcome. Victim's sister gets killed even if, for example, I do not shoot her. Consequently, my action is not wrong, according to the Parfittian Principle. (Note also that although the set containing my action, your action and Gertrude's action has an alternative whose outcome is intrinsically better, it also has a proper subset that has an alternative whose outcome is intrinsically better: For example, the set containing only Gertrude's action.)

4.4 The Non-Standard Version vs. the Parfittian Principle

In this section I shall argue that the Non-Standard Version is preferable to the Parfittian Principle. The reason I think so is that the Parfittian Principle has the wrong implications whereas the Non-Standard Version has the intuitively correct implications in Case Five:

Case Five

Case Five is exactly like Case Two except that my shot counterfactually depends on your shot. (I am simply the sort of person that wouldn't shoot someone unless someone else goes first.) We may imagine that you did shoot first, but my bullet nevertheless reached Victim before yours.

Intuitively, I act wrongly in this case. This case differs from Case Two – where it is intuitively clear that I act wrongly – only in that that my action counterfactually depends on yours. But that fact cannot really make my action right. However, The Parfittian Principle implies I do not act wrongly whereas the Non-Standard Version implies that I do. (Both principles obviously imply that you act wrongly. But that is as it should be.) In Chapter Two (section 3.3), I showed that my action is wrong in Case Five, according to the Non-Standard Version. So, it remains to be shown that my action is not wrong, according to the Parfittian Principle.

My action is a member of a set of actions in Case Five that has an alternative whose outcome is intrinsically better; namely, the set containing my action and your action. The outcome would have been intrinsically better if neither you nor I had shot victim. However, that set (i.e., the set containing your

action and my action) has a proper subset that has an alternative whose outcome is intrinsically better: i.e., the set containing your action. If you had not shot Victim then I would not have shot Victim and Victim would not then have died. So, my action is not wrong, according to the Parfitian Principle. (The set containing only my action does not have a better alternative. It would not have been better if I had not shot Victim since you would still have killed him.)

It is instructive to consider this case too (Parfit considers a relevantly similar case – he calls that case “Case Three” – and he seems to think that it is a virtue of his principle that it has the correct implication in that case):³⁶

Case Thirteen

Case Thirteen is exactly like Case Five except that by shooting I Victim save a further person from some bad pain. (We may assume that I shoot Victim because I know that I save this person from pain and that you would have killed Victim in any event.)

Note that the Non-Standard Version in this case implies – as it certainly should – that you, but not I, act wrongly. My action is not a redundant negative difference maker. Clause (1) is not satisfied since the outcome would have been intrinsically worse, if I had not shot Victim and you had shot him. The other person would then have suffered pain and you would nevertheless have killed Victim. Consequently, my action is not a negative redundant difference maker. This is true of your action as well. However, you have an alternative whose outcome is intrinsically better: The outcome would have intrinsically been better if you had not shot Victim. (Your action consequently satisfies clause (1) of the Non-Standard Version.) Granted, that other person would then have suffered that pain, but Victim wouldn’t have been killed (and we may assume that – in this case – this option is better).

³⁶ See Derek Parfit, *Reasons and Persons*, Oxford, Clarendon Press (1984), 71.

Chapter Six, Further Cases

1. Introduction

In the previous chapter I introduced the Non-Standard Version of Act-Consequentialism. I showed that this principle does have the intuitively correct implication in Case One (i.e., the over-determination case where you and I simultaneously shoot and kill a third person called “Victim”) and in a number of other important cases. However, there are also quite many other cases relevant in assessing the Non-Standard Version. One of them has occurred – i.e., Case Two – but there are also other cases. In this chapter, I shall try show that the Non-Standard Version has the intuitively correct implication in these cases too.

In section 2, I shall argue that the Non-Standard Version has the intuitively correct implications in Case Two (i.e., the pre-emption case where you and I shoot a third person called “Victim,” but where it is my bullet that actually kills him) and some other relevantly similar cases. In section 3, I try to show that the Non-Standard Version has the intuitively correct implications in cases like Case One but where the agents have alternatives whose outcome is either intrinsically better or intrinsically worse. In section 4, I argue that the Non-Standard Version has the intuitively correct implications in cases like Case One but where one of the events is a natural event.

2. Pre-Emption

In this section, I shall argue that the Non-Standard Version has the correct implications in Case Two (i.e., the pre-emption case where you and I shoot a third person called “Victim,” but where it is my bullet that actually kills him) and a number of other similar cases. In Case Two, you and I shoot at the same time. But what about cases where our actions are not simultaneous? I am going to argue that the Non-Standard Version has the intuitively correct implications in these cases too.

In section 2.1, I show that the Non-Standard Version has the intuitively correct implications in Case Two. In section 2.2, I discuss the worry that although it is plausible to say that I act wrongly in Case Two, it might not be plausible to say that you do so since you do not actually kill Victim. In Section

2.3, I show that the Non-Standard Version has the intuitively correct implications in cases that are similar to Case Two but where your action is performed before my action. In section 2.4, I show that the Non-Standard Version has the intuitively correct implications in cases that are similar to Case Two but where your action occurs after I have killed Victim. Finally, in section 2.5, I show that the Non-Standard Version has the intuitively correct implications in a case similar to Case Two, but where you do not actually shoot Victim.

2.1 Case Two

The Non-Standard Version has the intuitively correct implications in Case Two. Here is Case Two once more:

Case Two

You and I shoot one shot each at Victim. My bullet arrives first and kills Victim. Your bullet arrives very shortly after mine but does not kill Victim since my bullet has already done that. Your bullet would however have killed Victim, if my bullet had not done so first. Victim would not have been killed, if neither of us had pulled the trigger. Apart from shooting Victim, you and I have only one alternative each: each could also have abstained from shooting Victim. Furthermore, each of us would have performed his or her alternative – i.e., each of us would have abstained from shooting Victim – if he or she had not shot him. We act independently of one another, Victim’s level of well-being is not affected by my shot (e.g., Victim does not experience any pain), and no one else is affected. Finally, Victim would have gone on to live a life well worth living, if neither you nor I had shot him, and the extra time Victim would have lived, if I had not shot him would neither have been intrinsically good nor intrinsically bad for him.

Neither my action nor your action satisfies clause (1) of the Non-Standard Version: If I had not shot Victim you would still have done so, and the other way around. Consequently, if the Non-Standard Version has the intuitively correct implications in this case this must be because our actions satisfy clause (2) of the Non-Standard Version.

Both our actions are redundant negative difference makers also in Case Two, according to Redundant Negative Difference Making. Consider your action, for example. There is a combination – namely, my action – such that the clauses of Redundant Negative Difference Making are satisfied. Clause (1) is satisfied since you only have one alternative – i.e., your abstaining from shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse if you had abstained from shooting Victim and I had shot him since one shot is sufficient to kill him. Furthermore, clause (2) is satisfied since I only have one alternative – i.e., my abstaining from shooting Victim –

and the outcome would have been neither intrinsically better nor intrinsically worse if I had abstained from shooting Victim and you had shot Victim since (once more) one shot is sufficient to kill him. Finally, clause (3) is satisfied since we have an alternative – i.e., the combination “my abstaining from shooting Victim, and your abstaining from shooting Victim” – such that the outcome would have been intrinsically better, if we had performed that combination: Victim would under those circumstances not have been killed. (A parallel argument establishes that my action is a redundant negative difference maker in Case Two, according to Redundant Negative Difference Making.)

But do we have an alternative that would not be a redundant negative difference maker, according to Redundant Negative Difference Making? Yes, consider your alternative for example (i.e., your abstaining from shooting). Go to the closest possible world where you abstain from shooting. Clause (2) is not satisfied in that world since the outcome would be intrinsically better, if I had abstained from shooting Victim and you also had done so.

Parallel remarks show that my alternative would not be a negative redundant difference maker either, according to Redundant Negative Difference Making. So, the Non-Standard Version implies that you and I respectively act wrongly in Case Two.

2.2 A Worry

In Chapter Three, I noted that although it is intuitively clear that I act wrongly in Case Two, it is not as evident that you do so too since you do not actually kill Victim here. Yet, I suggested that it is not too implausible to say that you act wrongly in Case Two, and that a moral theory therefore may imply that you act wrongly here. After all, it is true that Victim is deprived of a future that would have been well worth living as a consequence of what we do, that each of our actions is causally sufficient for killing him, and that each of us would have killed him, even if the other had abstained from shooting. You are simply deeply involved in Victim’s bad fortune. In this section, I want to suggest that even if I am wrong about this (i.e., about it not being too implausible to say that you act wrongly in Case Two), the Non-Standard Version might still have some advantage over the Standard Version and the Causal Consequence Version.

Consider first the Standard Version. It has the intuitively wrong implications in Case Two since it implies that I do not act wrongly in Case Two. (As noted on several occasions, the Standard Version implies that neither you nor I act wrongly in Case Two since neither you nor I have an alternative whose outcome would be intrinsically better.) Consequently, both the Standard Version and the Non-Standard Version would then have the intuitively wrong implications in Case Two, whereas the Non-Standard Version but not the Standard Version would have the intuitively correct implications in Case One. (As also noted on several occasions, the Standard Version implies that neither you

nor I act wrongly in Case One too since neither you nor I have an alternative whose outcome would be intrinsically better here.) Now, since these two views otherwise have the same implications, the Non-Standard View has the advantage of getting one more type of case right.

Consider now the Causal Consequence Version. Provided that we accept that the state of affairs that Victim dies at t_1 is intrinsically bad for Victim, both the Non-Standard Version and the Causal Consequence Version would have the intuitively correct implications in Case One. The Causal Consequence would also have the intuitively correct implications in Case Two as opposed to the Non-Standard Version. However, we saw that the Causal Consequence Version has very implausible consequences in cases similar to Case One but where all my alternatives are such that they would harm some person (in a manner that is less morally serious than killing Victim). So, provided that we accept that the state of affairs that Victim dies at t_1 is intrinsically bad for Victim, these two views are roughly on par. But of course, we should probably not accept that the state of affairs that Victim dies at t_1 is intrinsically bad for Victim in the first place.

Before I move on, I want to note that if I am right about it not being too implausible to say that you act wrongly in Case Two (and that a moral theory therefore may imply that you act wrongly here), Case Two poses a very different challenge to moral theories than it does to theories of causation. A moral theory, I have suggested, *must* imply that I act wrongly, but *may* imply either that you do act wrongly or that you do not act wrongly. A theory of causation, on the other hand, *must* imply that I cause Victim's death and it *must* imply that you do not cause his death. If I am right about this, the moral theorist has some cause for celebration. It has proven extraordinary difficult to find an analysis of causation that implies that I but not you *cause* Victim's death.

2.3 When Your Action Occurs Before Mine

In the previous section I argued that we may accept a moral principle that implies that you and I, respectively, act wrongly in Case Two. Is this true of a relevantly similar case where your action occurs before mine? Consider the following case, for example:

Case Fourteen

You trick Victim into swallowing a very small bomb (in the shape of a pill), of a kind that stays in the body and without exception explodes and kills the person carrying it after an hour. Just before the bomb explodes, I stab Victim to death. Moreover, Victim would not have been killed, if neither of us had done what we did. Apart from doing what we did, you and I have only one alternative each: you could have abstained from tricking Victim

into swallowing a very small bomb, and I could have abstained from stabbing Victim. Furthermore, each of us would have performed his or her alternative if he or she had not done what he or she did. We act independently of one another, and no one else is affected. Finally, Victim would have gone on to live a life well worth living, if neither you nor I had done what we did, and the extra time Victim would have lived, if I had not stabbed him would neither have been intrinsically good nor intrinsically bad for him.

The Non-Standard Version implies that you and I act wrongly in this case too. Neither my action nor your action satisfies clause (1) of the Non-Standard Version: If I had stabbed Victim, he would still have been killed by the bomb, and if you had not tricked him into swallowing the bomb, I would still have stabbed him to death. However, both my action and your action satisfy clause (2) of the Non-Standard Version in Case Fourteen.

Both our actions are redundant negative difference makers in Case Fourteen, according to Redundant Negative Difference Making. Consider your action, for instance. There is a combination – namely, my action – such that the clauses of Redundant Negative Difference Making are satisfied. Clause (1) is satisfied since you only have one alternative – i.e., your abstaining from tricking Victim into swallowing a very small bomb – and the outcome would have been neither intrinsically better nor intrinsically worse if you had done that and I had stabbed him since my action is sufficient to kill Victim. Moreover, clause (2) is satisfied since I only have one alternative – i.e., my abstaining from stabbing Victim – and the outcome would have been neither intrinsically better nor intrinsically worse if I had abstained from stabbing him and you had tricked Victim into swallowing the bomb since the explosion is sufficient to kill him. Finally, clause (3) is satisfied since we have an alternative – i.e., the combination “my abstaining from stabbing Victim, and your abstaining from tricking Victim into swallowing a very small bomb” – such that the outcome would have been intrinsically better, if we had performed that combination: Victim would under those circumstances not have been killed.

A parallel argument establishes that my action is a redundant negative difference maker in Case Two, according to Redundant Negative Difference Making.

You and I, respectively, also have an alternative that would not be a redundant negative difference maker, according to Redundant Negative Difference Making. Consider, for example, my alternative. Go to the closest possible world where I perform my alternative. Clause (2) is not satisfied in that world since the outcome would be intrinsically better if you had abstained from tricking Victim into swallowing a bomb and I had abstained from stabbing him. Parallel remarks show that your alternative would not be a negative redundant

difference maker, according to Redundant Negative Difference Making. Consequently, the Non-Standard Version implies that you and I respectively act wrongly in Case Fourteen.

But is the implication that both you and I act wrongly in this case acceptable? Yes, I think so. It is intuitively clear that I act wrongly here. What about your action? In my view, it is clearer in this case that you act wrongly than it is in Case Two. It simply seems wrong to trick people into swallowing bombs of the kind in the example. Furthermore, as in Case One and Case Two, it is true that Victim is deprived of a future that would have been well worth living as a consequence of what we do, that each of our actions is causally sufficient for killing him, and that each of us would have killed him, even if the other had abstained from doing what he or she did. If I am right about this, the Non-Standard Version has the intuitively correct implication in Case Fourteen as well.

2.4 When Your Action Occurs After Mine

But what about a case like Case Two but where your action occurs after mine?

Consider the following case:

Case Fifteen

Victim is tied to a chair in a dark room. I enter the room and shoot and kill him, and leave. A couple of minutes later you enter the room and fire a bullet into Victim's dead body. You don't realise that Victim is dead. It looks perhaps like as if he is sleeping. Moreover, Victim would not have been killed, if neither of us had shot him. Apart from shooting Victim, you and I have only one alternative each: each could also have abstained from shooting Victim. Furthermore, each of us would have performed his or her alternative – i.e., each of us would have abstained from shooting Victim – if he or she had not shot him. We act independently of one another, and no one else is affected. Finally, Victim would have gone on to live a life well worth living, if neither you nor I had shot him, and the extra time Victim would have lived, if I had not shot him would neither have been intrinsically good nor intrinsically bad for him.

The Non-Standard Version implies that you and I act wrongly in Case Fifteen as well. Although neither satisfies clause (1) of the Non-Standard Version, both satisfy clause (2) of the Non-Standard Version.

My action is a redundant negative difference maker in this case, according to Redundant Negative Difference Making? There is a combination of actions – your action – such that the three clauses of Redundant Negative Difference Making are satisfied. Clause (1) is satisfied since I only have one alternative – i.e., my abstaining from shooting Victim – and the outcome would have been

neither intrinsically better nor intrinsically worse if I had performed that alternative and you had shot Victim since one shot is sufficient to kill him. Furthermore, clause (2) is satisfied since you only have one alternative – i.e., your abstaining from shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse if you had performed that alternative and I had shot him since (again) one shot is sufficient to kill him. Finally, clause (3) is satisfied since we have an alternative – i.e., the combination “my abstaining from shooting Victim, and your abstaining from shooting Victim” – such that the outcome would have been intrinsically better, if we had performed that combination: Victim would under those circumstances not have been killed.

An analogous argument establishes that your action is a redundant negative difference maker in Case Fifteen, according to Redundant Negative Difference Making.

But do we have an alternative that would not be a redundant negative difference maker, according to Redundant Negative Difference Making? Yes, consider your alternative for example (i.e., your abstaining from shooting). Go to the closest possible world where you perform your alternative. Clause (2) is not satisfied in that world since the outcome would be intrinsically better if I had abstained from shooting Victim and you also had done so. Parallel remarks show that my alternative would not be a negative redundant difference maker, according to Redundant Negative Difference Making. The Non-Standard Version consequently implies that you and I respectively act wrongly in Case Fifteen.

Is the implication that both you and I act wrongly in this case acceptable? Is it not obvious here that you do not act wrongly? Your action and the effect of your action both occur after the effect of my action. The deed is already a done deed when you enter the story. I nevertheless do not think that it is obvious that you do *not* act wrongly here. It is of course a lot less obvious that you *do* act wrongly in this case than it is in Case Fourteen. But (as I said) I do not think that is obvious that you don't. As in the other cases discussed in this section it is true that Victim is deprived of a future that would have been well worth living as a consequence of what we do, that each of our actions is causally sufficient for killing him, and that each of us would have killed him, even if the other had abstained from doing what he or she did.

2.5 When Your Action Does Not Occur

But what about a case like Case Two but where your action (i.e., the one sufficient for killing Victim) does not occur at all?

Here is a case like that:

Case Sixteen

I draw my gun, shoot and kill Victim. If I had not killed him, you would have done so a moment later. (You are a very quick shooter.) Seeing that I draw my gun, you remain motionless. Apart from doing what I did, I have only one alternative: I could only have abstained from shooting Victim, and that is also what I would have done, if I had not shot him. Apart from doing what you did, you could have shot Victim and you could have jumped on the spot. You would have performed one of your alternatives, if you had not remained motionless. I perform my alternative independent of you, but you would have shot Victim, if I had acted differently. Finally, no one else is affected and Victim lives a long and happy life in the closest possible world where neither you nor I shoot him.³⁷

Intuitively, my action is wrong in this case. However, it does not satisfy clause (1) of the Non-Standard Version: If I had not shot Victim you would still have done so. Furthermore, it is not evident that my action is a redundant negative difference maker, according to Redundant Negative Difference Making: You never fire your gun. It might therefore appear that the Non-Standard Version does not have the intuitively correct implications in Case Sixteen. However, we may focus on some previous action of yours. For example, we may focus on the decision you took some minutes before I shot Victim that you would draw your gun and shoot him, if you see that I do not do so. There is surely some such previous action (i.e., one that guarantees the truth of the counterfactual “If I had not shot him, you would have done so”). If that is the case, both my action and your action (i.e., your decision to shoot Victim if I do not), satisfy clause (2) of the Non-Standard Version.

My action is a redundant negative difference maker in Case Sixteen, according to Redundant Negative Difference Making. There is a combination of actions – your decision – such that the three clauses of Redundant Negative Difference Making are satisfied. Clause (1) is satisfied since I only have one alternative – i.e., my abstaining from shooting Victim – and the outcome would have been neither intrinsically better nor intrinsically worse if I had performed that alternative and you had taken your decision since you would then surely have killed him. Furthermore, clause (2) is satisfied since you only have one alternative – i.e., your abstaining from taking your decision – and the outcome would have been neither intrinsically better nor intrinsically worse if you had performed that alternative and I had shot him since (again) one shot is sufficient to kill him. Finally, clause (3) is satisfied since we have an alternative – i.e., the combination “my abstaining from shooting Victim,

³⁷ This case is like Case Six, i.e., a case of early pre-emption. See Chapter Six, section 4.4.

and your abstaining from taking your decision” – such that the outcome would have been intrinsically better, if we had performed that combination: Victim would under those circumstances not have been killed.

A parallel argument establishes that your decision is a redundant negative difference maker, according to Redundant Negative Difference Making.

But do we have an alternative that would not be a redundant negative difference maker, according to Redundant Negative Difference Making? Yes, consider your alternative for example (i.e., your refraining from taking the decision). Go to the closest possible world where you perform your alternative. Clause (2) is not satisfied in that world since the outcome would be intrinsically better if I had abstained from shooting Victim and you had refrained from taking the decision to shoot Victim if I do not do so.

Parallel remarks show that my alternative would not be a negative redundant difference maker, according to Redundant Negative Difference Making. So, the Non-Standard Version implies that you and I respectively act wrongly in Case Sixteen, after all.

3. Cases with Better or Worse Alternatives

The Non-Standard Version has the intuitively correct implications in Case One. An important feature of that case is that neither you nor I have an alternative whose outcome is intrinsically better. Another important feature of that case is that both you and I, respectively, have an alternative whose outcome is not intrinsically worse. What are the implications of the Non-Standard Version in cases that are as Case One but lack these features? In this section I am going to argue that the Non-Standard Version has the intuitively correct implications in these cases also.

In section 3.1, I consider two cases: One that is as Case One except that both you and I have an alternative whose outcome is intrinsically better, and one that is as Case One except that both you and I only have alternatives that are intrinsically worse. In section 3.2, I also consider two cases: One that is as Case One except that I but not you have an alternative whose outcome is intrinsically better, and one that is as Case One except that I but not you only have alternatives that are intrinsically worse. In section 3.3, I discuss and reject a modified version of the Non-Standard Version. (As we shall see, some might think that the Non-Standard Version has the wrong implications in the cases I discuss in section 3.2. The modified version of the Non-Standard Version has the implications they think are the correct implications in these cases.) In Section 3.4, I consider a case where you and I do not together produce the best outcome although each of us acts rightly. Finally, in section 3.5, I discuss some cases where the alternatives are only very slightly intrinsically worse but

where it would have been enormously much better if none of the agents involved had acted as they did. As I said, I argue that the Non-Standard Version has the correct implications in all these cases.

3.1 Cases Where All Agents Have Better or Worse Alternatives

In this section I argue that the Non-Standard Version has the correct implications both in a case that is as Case One except that both you and I have an alternative whose outcome is intrinsically better, but also in a case that is as Case One except that both you and I only have alternatives that are intrinsically worse.

Consider this case:

Case Seventeen

Case Seventeen is as Case One except that you and I, respectively, apart from Victim, kill ten further persons. (We do not kill the same ten persons.)

It is I think pretty evident that both you and I act wrongly in Case Seventeen. What does the Non-Standard Version say? It implies that both you and I act wrongly this case. It is true that neither of us satisfies clause (2) of the Non-Standard Version since neither of our actions is a redundant negative difference maker, according to Redundant Negative Difference Making. For example, neither of our actions satisfies clause (1) of Redundant Negative Difference Making since each of us has an alternative such that the outcome would have been intrinsically better if that alternative and the other person's action had been performed: Ten fewer people would then have been killed. However, this obviously implies that both of us satisfy clause (1) of the Non-Standard Version. Each of us consequently acts wrongly, according to the Non-Standard Version.

Consider now this case:

Case Eighteen

Case Eighteen is as Case One except that you and I have a different alternative this time. Apart from shooting Victim, you and I could only shoot ten other persons. (We would not have killed the same ten persons.)

Intuitively, neither you nor I act wrongly here. This is also what the Non-Standard Version implies. First, neither of our actions satisfies clause (2) of the Non-Standard Version since neither of our actions is a redundant negative difference maker, according to Redundant Negative Difference Making. For instance, neither of our actions satisfies clause (1) of Redundant Negative Difference Making since neither of us has an alternative such that the outcome

would not have been intrinsically worse if that alternative had been performed and the other person's action had been performed. If either of us had acted differently ten more people would have died. Second, neither you nor I satisfy clause (1) of the Non-Standard Version either since neither you nor I have an alternative whose outcome would be intrinsically better. Neither of us consequently acts wrongly in Case Eighteen, according to the Non-Standard Version.

3.2 Cases Where Some Agents Have Better or Worse Alternatives

In this section I argue that the Non-Standard Version has the correct implications both in a case that is as Case One except that I but not you have an alternative whose outcome is intrinsically better, and also in a case that is as Case One except that I but not you only have alternatives that are intrinsically worse.

Consider the following case:

Case Nineteen

This case is exactly as Case Seventeen except that you do not kill any further persons by what you do.

Now, the Non-Standard Version implies that I act wrongly but that you do not in Case Nineteen. First, neither of our actions is a redundant negative difference maker: My action is not a redundant negative difference maker, according to Redundant Negative Difference Making, since it does not satisfy clause (1) of Redundant Negative Difference Making: My action has an alternative such that the outcome would have been intrinsically better if that alternative – i.e., my abstaining from shooting Victim – had been performed and your shooting Victim had been performed. Your action is not a redundant negative difference maker, according to Redundant Negative Difference Making, since it does not satisfy clause (2) of Redundant Negative Difference Making. (It does not do so for the same reason that my action does not satisfy clause (1) of Redundant Negative Difference Making.) However, my action satisfies clause (1) of the Non-Standard Version whereas yours do not. I have an alternative whose outcome is intrinsically better, whereas you do not have such an alternative.

But would we not want to say that your action is wrong in Case Nineteen? In Case One, neither of us kills any further persons but both of us act wrongly, according to the Non-Standard Version. In the case under consideration – which is exactly like Case One except that I kill ten more people than I do in Case One – you suddenly do not act wrongly, according to the Non-Standard Version, although you do exactly what you do in Case One. But should really

the fact that I kill a lot of other persons – apart from also shooting Victim – affect the moral status of your action?

Well, I think that the implication that you do not act wrongly in Case Nineteen is acceptable. In Chapter Three, I noted that one reason why it seems plausible to say that you and I, respectively, act wrongly in Case One is that it seems pretty certain that something is not as it should be in Case One, and that other suggestions on what that might be did not seem very promising. For example, the suggestion that we act wrongly together did not seem very promising. Now, I did not consider the possibility that only one of us act wrongly in Case One since there is no relevant difference between what you do and I do in Case One. However, this is not the case here. Although it is clear that something is not as it should be, I also kill ten further persons here. This is why I think it is acceptable to say that I but not you act wrongly in this case while saying that both you and I act wrongly in Case One.

Consider now this case:

Case Twenty

This case is exactly as Case Eighteen except that you do not save any further persons by what you do.

The Non-Standard Version implies that neither you nor I act wrongly in Case Twenty too. Neither of our actions satisfy clause (1) of the Non-Standard Version since neither of us have an alternative whose outcome is intrinsically better. Moreover, neither of our actions satisfy clause (2) of the Non-Standard Version since neither of our actions is a redundant negative difference maker, according to Redundant Negative Difference Making.

My action is not a redundant negative difference maker, according to Redundant Negative Difference Making, since it does not satisfy clause (1) of Redundant Negative Difference Making: My action does not have an alternative such that the outcome would have been intrinsically better if that alternative – i.e., my abstaining from shooting Victim – had been performed, and your shooting Victim had been performed. Your action is not a redundant negative difference maker, according to Redundant Negative Difference Making, since it does not satisfy clause (2) of Redundant Negative Difference Making. (It does not do so for the same reason that my action does not satisfy clause (1) of Redundant Negative Difference Making.)

But do we really want to say that your action is right here? In Case One, neither of us kills any further persons but both of us act wrongly, according to the Non-Standard Version. In the case under consideration – which is exactly like Case One except that I also save ten people – you suddenly do not act wrongly, according to the Non-Standard Version, although you do exactly what you do in Case One. However, should the fact that I also save ten other persons affect the moral status of your action?

Well, just as I thought that the implication that you do not act wrongly in Case Nineteen is acceptable, I think that the implication that you do not act wrongly in Case Twenty is acceptable but for different reasons. In Case Eighteen, it is not clear that something is not as it should be. I obviously have an obligation to shoot Victim since I would otherwise kill those other ten persons. Therefore, I do not think we have to “find” someone to pin the blame on.

3.3 Modifying the Non-Standard Version

In this section I shall present a modified version of the Non-Standard Version that does imply that you act wrongly in Case Nineteen and Case Twenty. However, I do not think it is very attractive. As we shall see, it is quite complex, and we do not have to give up the Non-Standard Version in light of Case Nineteen and Case Twenty since (as I said), the implications that you do not act wrongly in these two cases are acceptable.

There are as far as I can see two things that are noteworthy as regards your action in Case Nineteen and in Case Twenty. First, your action is what I shall call a “person relative redundant negative difference maker.” (This is – as we shall see – true of my action as well.) Second, your action does not have an alternative that is intrinsically better, or one that is intrinsically worse. (This is not true of my action.)

I suppose the reader can form some idea (in view of our discussion in this and the previous chapter) what a person relative redundant negative difference maker is. But let me nevertheless spell this out:

Person Relative Redundant Negative Difference Making

An action a is a redundant negative difference maker relative to a person P , if and only if, there is a combination of actions C such that

(1) there is no alternative to a such that the outcome would be intrinsically better *for* P , if that alternative and the combination C had been performed, and there is at least one alternative to a such that the outcome would not be intrinsically worse *for* P if that alternative and the combination C had been performed,

(2) there is no alternative to the combination C such that the outcome would be intrinsically better *for* P , if that alternative and a had been performed, and there is at least one alternative to the combination C such that the outcome would not be intrinsically worse *for* P , if that alternative and a had been performed,

(3) there is an alternative to the combination a and C such that the outcome would be intrinsically better *for* P , if that alternative had been performed.

Both our actions in both Case Nineteen and Case Twenty are person relative redundant negative difference makers, according to Person Relative Redundant Negative Difference Making. It is enough if we consider your action in Case Nineteen. Clause (1) is satisfied since the outcome would have been neither intrinsically better nor intrinsically worse *for Victim* if you had abstained from shooting him and I had shot him since my action is sufficient to kill Victim. Furthermore, clause (2) is satisfied since the outcome would have been neither intrinsically better nor intrinsically worse *for Victim* if I had abstained from shooting him and you had shot him since one shot is sufficient to kill Victim. Finally, clause (3) is satisfied since the outcome would have been intrinsically better *for Victim* if neither you nor I had shot him. Your action is consequently a redundant negative difference maker relative Victim.

Let us modify the Non-Standard Version in view of this:

The Modified Non-Standard Version (its criterion of wrongness)

An action *a* is wrong, if and only if,

- (1) *a* has an alternative whose outcome would be intrinsically better than that of *a*, *or*
- (2)(a) *a* is a redundant negative difference maker, and
- (b) *a* has an alternative that would not be a redundant negative difference maker, *or*
- (3) (a) *a* is a redundant negative difference maker relative to some person, and
- (b) *a* has an alternative that would not be a redundant negative difference maker relative to some person, and
- (c) there is no alternative to *a* whose outcome is intrinsically better than that of *a*, and there is at least one alternative to *a* whose outcome is intrinsically worse than that of *a*.

Both our actions in Case Nineteen and Case Twenty are redundant negative difference makers relative to some person, according to Person Relative Redundant Negative Difference Making. However, neither of my two actions satisfies clause (3) of the Modified Non-Standard Version since neither satisfies clause (c) of that clause. In Case Nineteen, I have an alternative to what I do whose consequences would be intrinsically better, and in Case Twenty I have at least one alternative whose consequences would not be intrinsically worse. Both your actions clearly do, however. So, you act wrongly, according to the modified Non-Standard Version in both cases under consideration.

However, as I said, I do not think it is a good idea to modify the Non-Standard Principle in this manner. There is no big gain (since the Non-Standard Version does handle Case Nineteen and Case Twenty) and the modification is quite complex.

3.4 The Non-Standard Version is “Exclusively Act-Oriented”

Donald Regan has claimed that consequentialists are “inspired in the main by two distinct but equally fundamental intuitions.”³⁸

On the one hand, there is the intuition that whatever the correct moral theory is, it ought to be a good theory for *individuals* to follow as individuals. It ought to be the case that when an individual satisfies the theory, he produces the best consequences he can produce in the circumstances in which he finds himself. This is the intuition that underlies act-utilitarianism. On the other hand, there is the intuition that whatever the correct moral theory is, it ought to be a good theory for *everyone* to follow. It ought to be the case that if all agents satisfy the theory, then the class of all agents produce the best consequences they can produce collectively by any pattern of behaviour. This intuition gives rise to the varieties of rule-utilitarianism and utilitarian generalization.³⁹

The Non-Standard Version accommodates these intuitions in the cases we have studied up until now. It is in each case true of each agent that if the agent acts in accordance with the Non-Standard Version, then the agent produces the best outcome the agent can produce under the circumstances. (By saying that the agents act in accordance with the Non-Standard Version I mean that they do not act wrongly, according to the Non-Standard Version.) It is moreover in each case true that if each agent acts in accordance with the Non-Standard Version then the agents together produce the best outcome they together can produce under the circumstances. I can’t discuss all cases but it is (I take it) pretty evident that this is true. However, there are cases where this is not true.

Here is such a case:

Case Twenty-One

You and I shoot and kill Victim. Our bullets arrive at the same time and each shot would have killed Victim by itself. Furthermore, if I had shot Victim, but you had not, you would have killed ten other persons, and the other way around. Each of us only had one option apart from shooting Victim: to shoot in the direction of these people. However, if neither of us had shot Victim, no person would have get killed, not Victim nor those other people. We may imagine that our bullets under those circumstances would have clashed in mid-air and changed direction. We act independently of one another, Victim’s level of well-being is not affected by either of the shots taken by itself (e.g., Victim does not experience any pain), and no one

³⁸ Donald Regan, *Utilitarianism and Co-operation*, Clarendon Press, Oxford, 1980, 3.

³⁹ *Ibid.* 3.

else is affected. Finally, Victim would have gone on to live a life well worth living, if neither you nor I had shot him.⁴⁰

Note that neither of us acts wrongly, according to the Non-Standard Version. Neither of our actions satisfies clause (1) of the Non-Standard Version since neither of us has an alternative whose outcome would be intrinsically better. Furthermore, neither of us satisfies clause (2) of the Non-Standard Version since neither of our actions is a redundant negative difference maker, according to Redundant Negative Difference Making. For example, neither of our actions satisfies clause (1) of Redundant Negative Difference Making. This is so because neither of us has an alternative such that the outcome would not have been intrinsically worse if that alternative had been performed and the other person's action had been performed. If either of us had acted differently ten more people would have died. Note, however, that it is not true in this case that you and I together produce the best outcome we (together) can produce under the circumstances: It would have been better if neither of us had shot Victim.

Now, it may be argued that this case is a reason to reject the Non-Standard Version. It is – it might be claimed – pretty pointless to adopt a principle (like the Non-Standard Version) that considers combinations of actions but fails to handle Case Twenty-One (and other similar cases). If there is a case such principle should be able to handle, this is the one, it might be claimed. Well, I think the Non-Standard Version does handle this case. Provided that you shoot Victim, that is what I should do too, and the other way around. There is simply no way around this difficulty (i.e., that you and I, respectively, do not act wrongly although it would have been better if neither of us had shot Victim). Donald Regan has claimed that no moral theory that is “exclusively act-oriented” handles cases like Case Twenty-One. A moral theory is roughly, according to Regan, “exclusively act-oriented if it can be stated in the form ‘An agent should do that act which...’”⁴¹ So, insofar that we want our theory to be exclusively act-oriented (which we do) the Non-Standard Version is as good as it gets.⁴²

⁴⁰ This case is similar to a case presented by Donald Regan. *Ibid.* 18.

⁴¹ *Ibid.* 10.

⁴² In Chapter 3, I said that there are cases similar to Case One and Case Two where Rule-Consequentialism (which claims that action is wrong if it does not conform to an ideal code of rules) and Consequentialist Generalisation (which claims that an action is wrong, if it is not the case that the outcome would be best if everyone did it). Case Twenty-One is such a case: These views imply that you and I, respectively, act wrongly by shooting Victim. But this is counter-intuitive.

3.5 A Further Objection

Recall that the Non-Standard Version handles cases like Case Eighteen where the agents involved only have alternatives whose consequences are intrinsically worse. However, what about cases where the alternatives are only very slightly intrinsically worse but where it would have been enormously much better if none of the agents involved had acted as they did?

Here is a case like that:

Case Twenty-Two

There are two buttons. If either is pressed 1.000 children die in extreme agony. If neither is pressed no child dies. Moreover, if the left button is pressed, a funny noise is produced that will make you laugh uncontrollably. If the right button is pressed a different funny noise is produced that will make me laugh uncontrollably. I press the left button and you press the right one. You and I have the time of our lives (we press the buttons again and again to roars of laughter); and 1.000 children die in extreme agony. Neither you nor I affect anyone else.

Neither you nor I act wrongly in this case, according to the Non-Standard Version. Neither of our actions satisfies clause (1) of the Non-Standard Version since neither of us has an alternative whose consequences are intrinsically better. Neither of us satisfies clause (2) of the Non-Standard Version either since neither of our actions is a redundant negative difference maker: neither of us has an alternative such that the outcome would not be intrinsically worse if that alternative and the other person's action had been performed. If I hadn't pressed the button, you wouldn't have had the time of your life, and the other way around. However, these implications might seem counterintuitive.

The reason (I suspect) that one might think that, for example, I act wrongly, in Case Twenty-Two, is that the gain that is produced by me pressing the button is small compared to the loss produced by you and me pressing the button. However, I do not think that we should accommodate this intuition. As far as I can see, there is no principled way to balance the gain produced by the individual action against the loss produced by the compound action. Consider the following case, for example:

Case Twenty-Three

There are two buttons. If either is pressed the same 1.000 children die in extreme agony. If neither is pressed no child dies. Moreover, if the left button is pressed, 300 children are saved. If the right button is pressed another 300 children are saved. I press the left button and you press the right one. 1.000 children die in extreme agony; but 600 children are also saved. No one else is affected.

In this case, the gain that is produced by me pressing the button is also comparatively small (although it is not as small as in the previous case). However, I clearly ought to press the button in this case. Otherwise 300 children won't be saved. Although this example is by no means decisive, I nevertheless think that it illustrates the difficulty confronting anyone who wants to accommodate the intuition that you and I act wrongly in Case Twenty-Two. I shall therefore suggest that we leave the Non-Standard Version as it is.

4. Cases with Natural Causes

In the cases we have studied thus far the relevant events have been actions. What should we say in cases where this is not so? Consider for example the following case:

Case Twenty-Four

I shoot Victim at the same time as a falling rock hits his head. To be more precise, my bullet pierces Victim's heart at the same time as the rock crushes his head. Victim naturally dies as a result of this. My shot would have killed Victim by itself, and the falling rock would have killed Victim by itself. However, Victim would not have been killed, if neither of these two events had happened. Furthermore, no one else is affected, and Victim would have gone on to live a life well worth living, if I had not shot him and if the rock had not hit him. Finally, my action and the falling rock are independent of one another and I had one alternative apart from shooting Victim: I could have abstained from shooting him, and this is also what I would have done, if I had not shot Victim.

The Non-Standard Version implies that I do not act wrongly in this case. My action does not satisfy clause (1) of the Non-Standard Version: The outcome would not have been intrinsically better if I had not shot Victim since he would at any rate have been killed by the falling rock. My action does not satisfy clause (2) of the Non-Standard Version either: Redundant Negative Difference Making is concerned with actions, but one of the events in this case is not an action, namely the falling rock.

But should we not say that I do act wrongly in Case Twenty-Four? Well, I think we may say that I do not act wrongly here. There is an important difference between this case and for example Case One. There is no action nor a combination of actions that is such that it has an alternative whose outcome is intrinsically better in Case Twenty-Four, whereas there is such a combination

in Case One (i.e., the combination, “my abstaining from shooting Victim, your abstaining from shooting Victim”). The outcome is in this sense inevitable.)⁴³

⁴³ It may be interesting to observe that the law holds, in cases of over-determination where one of the causes is an action and the other is a natural event, that the actor is not liable. See Michael Moore on what the law has to say as regards various cases of redundant causation. Michael S. Moore, *Causation and Responsibility*, New York: Oxford University Press (2009), chap. 4–6 (see also 115).

Bibliography

- Bennett, J. (1988) *Events and their Names*, Indianapolis, Hackett Publishers.
- Berker, S. (2009) "The Normative Insignificance of Neuroscience," *Philosophy & Public Affairs* 37: 293–329.
- Bratman, M.E. (2014) *Shared Agency: A Planning Theory of Acting Together*, Oxford University Press.
- Carlson, E. (1995) *Consequentialism Reconsidered*, Dordrecht, Kluwer Academic Publishers.
- Goodman, A. (1970) *A Theory of Human Action*, Princeton University Press.
- Jackson, F. (1987) "Group Morality," *Metaphysics and Morality: Essays in Honour of J.J. C. Smart*, ed. Pettit, P., Sylvan, R., and Norman, J., Oxford: Basil Blackwell, 91–110.
- Jackson, F. (1997) "Which Effects?," *Reading Parfit*, Blackwell, 42–53.
- Kagan, S. (2011) "Do I Make a Difference?" *Philosophy & Public Affairs* 39: 105–140.
- Lewis, D. (1973) "Causation," *Journal of Philosophy* 70: 556–567.
- Lewis, D. (1986) "Postscripts to 'Causation'," in David Lewis *Philosophical Papers Vol. II*, Oxford, Oxford University press, 172–213.
- Lewis, D. (2004) "Causation as Influence," in Collins, J., Hall N., Paul L. A. (ed.), *Causation and Counterfactuals*, Cambridge Massachusetts, MIT Press, 74–106.
- List, C. and Pettit, P. (2011) *The Possibility, Design, and Status of Corporate Agents*, Oxford University Press.
- Moore, M. S. (2009) *Causation and Responsibility*, New York: Oxford University Press.

- Norcross, A. (2005) "Harming in Context," *Philosophical Studies* 123: 149–173.
- Parfit, D. (1984) *Reasons and Persons*, Oxford, Clarendon Press.
- Pinkert, F. (2015) "What If I Cannot Make a Difference (and Know it)?" *Ethics* 125: 971–998.
- Regan, D. (1980) *Utilitarianism and Co-operation*, Clarendon Press, Oxford.
- Singer, P. (2005) "Ethics and Intuitions," *Journal of Ethics* 9: 331–352.
- Sumner, L. W. (1999) *Welfare, Happiness, and Ethics*, Oxford University Press.
- Tersman, F. (2008) "The reliability of moral intuitions: A challenge from neuroscience," *Australasian Journal of Philosophy* 86: 389–405.
- Timmons, M. (2013) *Moral Theory: An Introduction*, Second Edition, Rowman & Littlefield Publishers, Inc.
- Tännsjö, T. (1989) "The Morality of Collective Actions", *The Philosophical Quarterly* 39: 221–228.
- Zimmerman, M. J. (1996) *The Concept of Moral Obligation*, Cambridge: Cambridge University Press.

